

GLOBAL
EDITION



Database Processing

Fundamentals, Design, and Implementation

FOURTEENTH EDITION

David M. Kroenke • David J. Auer

ALWAYS LEARNING

PEARSON

OTHER MIS TITLES OF INTEREST

Introductory MIS

Managing Information Technology, 7/e
Brown, DeHayes, Hoffer, Martin & Perkins ©2012

Experiencing MIS, 6/e
Kroenke & Boyle ©2016

Using MIS, 8/e
Kroenke & Boyle ©2016

MIS Essentials, 4/e
Kroenke ©2015

Management Information Systems, 14/e
Laudon & Laudon ©2016

Essentials of Management Information Systems, 11/e
Laudon & Laudon ©2015

IT Strategy, 3/e
McKeen & Smith ©2015

**Processes, Systems, and Information:
An Introduction to MIS, 2/e**
McKinney & Kroenke ©2015

Information Systems Today, 7/e
Valacich & Schneider ©2016

Introduction to Information Systems, 2/e
Wallace ©2015

Database

Hands-on Database, 2/e
Conger ©2014

Modern Database Management, 12/e
Hoffer, Ramesh & Topi ©2016

**Database Systems: Introduction to Databases
and Data Warehouses**
Jukic, Vrbsky & Nestorov ©2014

Essentials of Database Management
Hoffer, Topi & Ramesh ©2014

Database Concepts, 7/e
Kroenke & Auer ©2015

Database Processing, 14/e
Kroenke & Auer ©2016

Systems Analysis and Design

Modern Systems Analysis and Design, 7/e
Hoffer, George & Valacich ©2014

Systems Analysis and Design, 9/e
Kendall & Kendall ©2014

Essentials of Systems Analysis and Design, 6/e
Valacich, George & Hoffer ©2015

Decision Support Systems

Business Intelligence, 3/e
Sharda, Delen & Turban ©2014

**Decision Support and Business
Intelligence Systems, 10/e**
Sharda, Delen & Turban ©2014

Data Communications & Networking

Applied Networking Labs, 2/e
Boyle ©2014

Digital Business Networks
Dooley ©2014

Business Data Networks and Security, 10/e
Panko & Panko ©2015

Electronic Commerce

**E-Commerce: Business, Technology,
Society, 11/e**
Laudon & Traver ©2015

Enterprise Resource Planning

Enterprise Systems for Management, 2/e
Motiwalla & Thompson ©2012

Project Management

**Project Management: Process,
Technology and Practice**
Vaidyanathan ©2013



Database Processing

Fundamentals, Design,
and Implementation

14TH EDITION
Global Edition



This page intentionally left blank



Database Processing

Fundamentals, Design,
and Implementation


14TH EDITION
Global Edition



David M. Kroenke

David J. Auer

Western Washington University



PEARSON

Boston Columbus Indianapolis New York San Francisco Hoboken
Amsterdam Cape Town Dubai London Madrid Milan Munich Paris Montréal Toronto
Delhi Mexico City São Paulo Sydney Hong Kong Seoul Singapore Taipei Tokyo

Vice President, Business Publishing: Donna Battista
Editor in Chief: Stephanie Wall
Acquisitions Editor: Nicole Sam
Program Manager Team Lead: Ashley Santora
Program Manager: Denise Weiss
Editorial Assistant: Olivia Vignone
Vice President, Product Marketing: Maggie Moylan
Director of Marketing, Digital Services and Products:
Jeanette Koskinas
Executive Product Marketing Manager: Anne Fahlgren
Field Marketing Manager: Lenny Ann Raper
Senior Strategic Marketing Manager: Erin Gardner
Product Marketing Assistant: Jessica Quazza
Project Manager Team Lead: Jeff Holcomb
Project Manager: Ilene Kahn
Operations Specialist: Diane Peirano
Senior Art Director: Janet Slowik
Text Designer: Integra Software Services Pvt. Ltd.
Cover Designer: Lumina Datamatics, Inc.

Cover Photo: Omelchenko/Shutterstock
Vice President, Director of Digital Strategy & Assessment: Paul Gentile
Manager of Learning Applications: Paul Deluca
Digital Editor: Brian Surette
Digital Studio Manager: Diane Lombardo
Digital Studio Project Manager: Robin Lazrus
Digital Studio Project Manager: Alana Coles
Digital Studio Project Manager: Monique Lawrence
Digital Studio Project Manager: Regina DaSilva
Senior Manufacturing Controller, Global Edition: Trudy Kimber
Manager, Media Production, Global Edition: M. Vikram Kumar
Acquisitions Editor, Global Edition: Steven Jackson
Associate Project Editor, Global Edition: Priyanka Shivas
Full-Service Project Management and Composition: Integra Software Services Pvt. Ltd.
Printer/Binder: Vivar, Malaysia
Cover Printer: Vivar, Malaysia
Text Font: 10/12 Mentor Std Light

Credits and acknowledgments borrowed from other sources and reproduced, with permission, in this textbook appear on the appropriate page within text.

Microsoft and/or its respective suppliers make no representations about the suitability of the information contained in the documents and related graphics published as part of the services for any purpose. All such documents and related graphics are provided "as is" without warranty of any kind. Microsoft and/or its respective suppliers hereby disclaim all warranties and conditions with regard to this information, including all warranties and conditions of merchantability, whether express, implied or statutory, fitness for a particular purpose, title and non-infringement. In no event shall Microsoft and/or its respective suppliers be liable for any special, indirect or consequential damages or any damages whatsoever resulting from loss of use, data or profits, whether in an action of contract, negligence or other tortious action, arising out of or in connection with the use or performance of information available from the services.

The documents and related graphics contained herein could include technical inaccuracies or typographical errors. Changes are periodically added to the information herein. Microsoft and/or its respective suppliers may make improvements and/or changes in the product(s) and/or the program(s) described herein at any time. Partial screen shots may be viewed in full within the software version specified.

Microsoft® Windows®, and Microsoft Office® are registered trademarks of the Microsoft Corporation in the U.S.A. and other countries. This book is not sponsored or endorsed by or affiliated with the Microsoft Corporation.

MySQL®, the MySQL Command Line Client®, the MySQL Workbench®, and the MySQL Connector/ODBC® are registered trademarks of Sun Microsystems, Inc./Oracle Corporation. Screenshots and icons reprinted with permission of Oracle Corporation. This book is not sponsored or endorsed by or affiliated with Oracle Corporation.

Oracle Database 12c and Oracle Database Express Edition 11g Release 2 2014 by Oracle Corporation. Reprinted with permission. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Mozilla 35.104 and Mozilla are registered trademarks of the Mozilla Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

PHP is copyright The PHP Group 1999–2012, and is used under the terms of the PHP Public License v3.01 available at http://www.php.net/license/3_01.txt. This book is not sponsored or endorsed by or affiliated with The PHP Group.

Pearson Education Limited
Edinburgh Gate
Harlow
Essex CM20 2JE
England

and Associated Companies throughout the world

Visit us on the World Wide Web at:
www.pearsonglobaleditions.com

© Pearson Education Limited 2016

The rights of David M. Kroenke and David J. Auer to be identified as the authors of this work has been asserted by them in accordance with the Copyright, Designs and Patents Act 1988.

Authorized adaptation from the United States edition, entitled Database Processing: Fundamentals, Design, and Implementation, 14/e, ISBN 978-0-13-387670-3, by David M. Kroenke and David J. Auer, published by Pearson Education © 2016.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without either the prior written permission of the publisher or a license permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd, Saffron House, 6–10 Kirby Street, London EC1N 8TS.

All trademarks used herein are the property of their respective owners. The use of any trademark in this text does not vest in the author or publisher any trademark ownership rights in such trademarks, nor does the use of such trademarks imply any affiliation with or endorsement of this book by such owners.

ISBN 10: 1-292-10763-4
ISBN 13: 978-1-292-10763-9

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

14 13 12 11 10 9 8 7 6 5 4 3 2 1

Typeset by Integra Software Services Pvt. Ltd. in Mentor Std Light, 10/12 pt.
Printed and bound in Malaysia

Brief Contents

PART 1 ■ Getting Started

33

- Chapter 1 Introduction 34
- Chapter 2 Introduction to Structured Query Language 68

PART 2 ■ Database Design

165

- Chapter 3 The Relational Model and Normalization 166
- Chapter 4 Database Design Using Normalization 209
- Chapter 5 Data Modeling with the Entity-Relationship Model 228
- Chapter 6 Transforming Data Models into Database Designs 280

PART 3 ■ Database Implementation

333

- Chapter 7 SQL for Database Construction and Application Processing 334
- Chapter 8 Database Redesign 428

PART 4 ■ Multiuser Database Processing

455

- Chapter 9 Managing Multiuser Databases 456
- Chapter 10 Managing Databases with Microsoft SQL Server 2014, Oracle Database, and MySQL 5.6 490
- Online Chapter: See Page 495 for Instructions**
- Chapter 10A Managing Databases with Microsoft SQL Server 2014
- Online Chapter: See Page 495 for Instructions**
- Chapter 10B Managing Databases with Oracle Database
- Online Chapter: See Page 495 for Instructions**
- Chapter 10C Managing Databases with MySQL 5.6

PART 5 ■ Database Access Standards

497

- Chapter 11 The Web Server Environment 498
- Chapter 12 Big Data, Data Warehouses, and Business Intelligence Systems 565
- Online Appendices: See Page 610 for Instructions**
- Appendix A Getting Started with Microsoft Access 2013
- Appendix B Getting Started with Systems Analysis and Design
- Appendix C E-R Diagrams and the IDEF1X Standard
- Appendix D E-R Diagrams and the UML Standard
- Appendix E Getting Started with the MySQL Workbench Data Modeling Tools
- Appendix F Getting Started with Microsoft Visio 2013
- Appendix G Data Structures for Database Processing
- Appendix H The Semantic Object Model
- Appendix I Getting Started with Web Servers, PHP, and the NetBeans IDE
- Appendix J Business Intelligence Systems
- Appendix K Big Data

This page intentionally left blank

Preface 19

PART 1 ■ Getting Started

33

Chapter 1: Introduction 34

- Chapter Objectives 34
- The Importance of Databases in the Internet and Smartphone World 35
- The Characteristics of Databases 37
 - A Note on Naming Conventions* 38 • *A Database Has Data and Relationships* 39 • *Databases Create Information* 40
- Database Examples 41
 - Single-User Database Applications* 41 • *Multuser Database Applications* 41 • *E-Commerce Database Applications* 42 • *Reporting and Data Mining Database Applications* 43
- The Components of a Database System 43
 - Database Applications and SQL* 44 • *The DBMS* 46 • *The Database* 47
- Personal Versus Enterprise-Class Database Systems 49
 - What Is Microsoft Access?* 49 • *What Is an Enterprise-Class Database System?* 50
- Database Design 52
 - Database Design from Existing Data* 52 • *Database Design for New Systems Development* 54 • *Database Redesign* 54
- What You Need to Learn 55
- A Brief History of Database Processing 56
 - The Early Years* 56 • *The Emergence and Dominance of the Relational Model* 58
 - *Post-Relational Developments* 59
- Summary 61 • Key Terms 62 • Review Questions 63 • Project Questions 65**

Chapter 2: Introduction to Structured Query Language 68

- Chapter Objectives 68
- Cape Codd Outdoor Sports 69
- Business Intelligence Systems and Data Warehouses 70
 - The Cape Codd Outdoor Sports Extracted Retail Sales Data* 71 • *RETAIL_ORDER Data* 72
 - *ORDER_ITEM Data* 74 • *SKU_DATA Table* 74 • *CATALOG_SKU_20### Tables* 75
 - *The Complete Cape Codd Data Extract Schema* 75 • *Data Extracts Are Common* 76
- SQL Background 76
- The SQL SELECT/FROM/WHERE Framework 77
 - Reading Specified Columns from a Single Table* 78 • *Specifying Column Order in SQL Queries from a Single Table* 79
- Submitting SQL Statements to the DBMS 80
 - Using SQL in Microsoft Access 2013* 80 • *Using SQL in Microsoft SQL Server 2014* 85 • *Using SQL in Oracle Database* 88 • *Using SQL in Oracle MySQL 5.6* 90
- SQL Enhancements for Querying a Single Table 93
 - Reading Specified Rows from a Single Table* 93 • *Reading Specified Columns and Rows from a Single Table* 97 • *Sorting the SQL Query Results* 97 • *SQL WHERE Clause Options* 100

Performing Calculations in SQL Queries	107
Using SQL Built-in Aggregate Functions	107
SQL Expressions in SQL SELECT Statements	111
Grouping Rows in SQL SELECT Statements	114
Querying Two or More Tables with SQL	119
Querying Multiple Tables with Subqueries	119
Querying Multiple Tables with Joins	122
Comparing Subqueries and Joins	127
The SQL JOIN ON Syntax	127
Outer Joins	130
Using SQL Set Operators	134
Summary	137
Key Terms	138
Review Questions	139
Project Questions	146
Case Questions	149
The Queen Anne Curiosity Shop	153
Morgan Importing	161

PART 2 ■ Database Design

165

Chapter 3: The Relational Model and Normalization 166

Chapter Objectives	166
Relational Model Terminology	168
Relations	168
Characteristics of Relations	169
Alternative Terminology	171
To Key, or Not to Key—That Is the Question!	172
Functional Dependencies	172
Finding Functional Dependencies	174
Keys	177
Normal Forms	180
Modification Anomalies	180
A Short History of Normal Forms	181
Normalization Categories	182
From First Normal Form to Boyce-Codd Normal Form Step by Step	182
Eliminating Anomalies from Functional Dependencies with BCNF	187
Eliminating Anomalies from Multivalued Dependencies	196
Fifth Normal Form	199
Domain/Key Normal Form	199
Summary	200
Key Terms	200
Review Questions	201
Project Questions	203
Case Questions	204
The Queen Anne Curiosity Shop	205
Morgan Importing	207

Chapter 4: Database Design Using Normalization 209

Chapter Objectives	209
Assess Table Structure	210
Designing Updatable Databases	211
Advantages and Disadvantages of Normalization	211
Functional Dependencies	212
Normalizing with SQL	212
Choosing Not to Use BCNF	213
Multivalued Dependencies	214
Designing Read-Only Databases	214
Denormalization	215
Customized Duplicated Tables	215
Common Design Problems	217
The Multivalued, Multicolumn Problem	218
Inconsistent Values	219
Missing Values	220
The General-Purpose Remarks Column	221
Summary	222
Key Terms	222
Review Questions	223
Project Questions	225
Case Questions	225
The Queen Anne Curiosity Shop	226
Morgan Importing	227

Chapter 5: Data Modeling with the Entity-Relationship Model 228

Chapter Objectives	228
The Purpose of a Data Model	229
The Entity-Relationship Model	229
Entities	229
Attributes	230
Identifiers	230
Relationships	231
Maximum Cardinality	233
Minimum Cardinality	234
Entity-Relationship Diagrams and Their Versions	235
Variations of the E-R Model	235
E-R Diagrams Using the	

- IE Crow's Foot Model 236 • Strong Entities and Weak Entities 238 • ID-Dependent Entities 238 • Non-ID-Dependent Weak Entities 239 • The Ambiguity of the Weak Entity 240 • Subtype Entities 240
- Patterns in Forms, Reports, and E-R Models 243
 - Strong Entity Patterns 243 • ID-Dependent Relationships 247 • Mixed Identifying and Nonidentifying Patterns 253 • The For-Use-By Pattern 256 • Recursive Patterns 257
- The Data Modeling Process 260
 - The College Report 261 • The Department Report 261 • The Department/Major Report 262 • The Student Acceptance Letter 264
- Summary 266 • Key Terms 267 • Review Questions 268 • Project Questions 270 • Case Questions 276 • The Queen Anne Curiosity Shop 278 • Morgan Importing 279**

Chapter 6: Transforming Data Models into Database Designs 280

- Chapter Objectives 280
- The Purpose of a Database Design 281
- Create a Table for Each Entity 281
 - Selecting the Primary Key 281 • Specifying Alternate Keys 284 • Specifying Column Properties 284 • Verify Normalization 291
- Create Relationships 292
 - Relationships Between Strong Entities 292 • Relationships Using ID-Dependent Entities 295 • Relationships with a Weak Non-ID-Dependent Entity 299 • Relationships in Mixed Entity Designs 300 • Relationships Between Supertype and Subtype Entities 302 • Recursive Relationships 303 • Representing Ternary and Higher-Order Relationships 304 • Relational Representation of the Highline University Data Model 306
- Design for Minimum Cardinality 309
 - Actions when the Parent Is Required 310 • Actions when the Child Is Required 311 • Implementing Actions for M-O Relationships 312 • Implementing Actions for O-M Relationships 313 • Implementing Actions for M-M Relationships 313 • Designing Special Case M-M Relationships 314 • Documenting the Minimum Cardinality Design 314 • An Additional Complication 315 • Summary of Minimum Cardinality Design 316
- The View Ridge Gallery Database 316
 - View Ridge Gallery Database Summary of Requirements 317 • The View Ridge Data Model 318 • Database Design with Data Keys 319 • Minimum Cardinality Enforcement for Required Parents 320 • Minimum Cardinality Enforcement for the Required Child 321 • Column Properties for the View Ridge Database Design Tables 323
- Summary 325 • Key Terms 327 • Review Questions 327 • Project Questions 329 • Case Questions 330 • The Queen Anne Curiosity Shop 332 • Morgan Importing 332**

PART 3 ■ Database Implementation

333

Chapter 7: SQL for Database Construction and Application Processing 334

- Chapter Objectives 334
- The Importance of Working with an Installed DBMS Product 335
- The View Ridge Gallery Database 335
- SQL DDL and DML 335
- Managing Table Structure with SQL DDL 337
 - Creating the VRG Database 337 • Using SQL Scripts 337 • Using the SQL CREATE TABLE Statement 338 • Variations in SQL Data Types and SQL/PSM 339 • Creating the VRG Database ARTIST Table 339 • Creating the VRG Database WORK Table and the 1:N ARTIST-to-WORK Relationship 342 • Implementing Required Parent Rows 343 • Implementing 1:1 Relationships 344 • Casual Relationships 344 • Creating Default Values and Data

- Constraints with SQL 344 • Creating the VRG Database Tables 346 • The SQL ALTER TABLE Statement 349 • The SQL DROP TABLE Statement 350 • The SQL TRUNCATE TABLE Statement 351 • The SQL CREATE INDEX Statement 351
- SQL DML Statements 352
 - The SQL INSERT Statement 352 • Populating the VRG Database Tables 353 • The SQL UPDATE Statement 359 • The SQL MERGE Statement 360 • The SQL DELETE Statement 361
- Using SQL Views 361
 - Using SQL Views to Hide Columns and Rows 364 • Using SQL Views to Display Results of Computed Columns 366 • Using SQL Views to Hide Complicated SQL Syntax 366 • Layering Built-in Functions 367 • Using SQL Views for Isolation, Multiple Permissions, and Multiple Triggers 369 • Updating SQL Views 370
- Embedding SQL in Program Code 371
 - SQL/Persistent Stored Modules (SQL/PSM) 373 • Using SQL User-Defined Functions 373 • Using SQL Triggers 377 • Using Stored Procedures 382 • Comparing User-Defined Functions, Triggers, and Stored Procedures 386
- Summary 386 • Key Terms 388 • Review Questions 389 • Project Questions 398 • Case Questions 401 • The Queen Anne Curiosity Shop 415 • Morgan Importing 422**

Chapter 8: Database Redesign 428

- Chapter Objectives 428
- The Need for Database Redesign 429
- SQL Statements for Checking Functional Dependencies 429
 - What Is a Correlated Subquery? 430
- How Do I Analyze an Existing Database? 435
 - Reverse Engineering 436 • Dependency Graphs 437 • Database Backup and Test Databases 437
- Changing Table Names and Table Columns 438
 - Changing Table Names 438 • Adding and Dropping Columns 440 • Changing a Column Data Type or Column Constraints 441 • Adding and Dropping Constraints 442
- Changing Relationship Cardinalities 442
 - Changing Minimum Cardinalities 442 • Changing Maximum Cardinalities 443
- Adding and Deleting Tables and Relationships 446
- Forward Engineering 446
- Summary 447 • Key Terms 448 • Review Questions 448 • Project Questions 450 • Case Questions 451 • The Queen Anne Curiosity Shop 452 • Morgan Importing 453**

PART 4 ■ Multiuser Database Processing

455

Chapter 9: Managing Multiuser Databases 456

- Chapter Objectives 456
- The Importance of Working with an Installed DBMS Product 457
- Database Administration 457
 - Managing the Database Structure 458
- Concurrency Control 459
 - The Need for Atomic Transactions 460 • Resource Locking 463 • Optimistic Versus Pessimistic Locking 465 • SQL Transaction Control Language and Declaring Lock Characteristics 466 • Implicit and Explicit COMMIT TRANSACTION 467 • Consistent Transactions 468 • Transaction Isolation Level 469 • SQL Cursors 470
- Database Security 472
 - Processing Rights and Responsibilities 472 • DBMS Security 473 • DBMS Security Guidelines 474 • Application Security 475 • The SQL Injection Attack 476

Database Backup and Recovery 477

Recovery via Reprocessing 477 • *Recovery via Rollback/Rollforward* 478

Managing the DBMS 480

Maintaining the Data Repository 481

**Summary 482 • Key Terms 483 • Review Questions 484 • Project Questions 485
• Case Questions 486 • The Queen Anne Curiosity Shop 487 • Morgan
Importing 488**

Chapter 10: Managing Databases with Microsoft SQL Server 2014, Oracle Database, and MySQL 5.6 490

Chapter Objectives 490

Installing the DBMS 491

Using the DBMS Database Administration and Database Development Utilities 492

Creating a Database 492

Creating and Running SQL Scripts 492

Reviewing the Database Structure in the DBMS GUI Utility 493

Creating and Populating the View Ridge Gallery VRG Database Tables 493

Creating SQL Views for the View Ridge Gallery VRG Database 493

Database Application Logic and SQL/Persistent Stored Modules (SQL/PSM) 493

DBMS Concurrency Control 494

DBMS Security 494

DBMS Database Backup and Recovery 494

Other DBMS Topics Not Discussed 494

Choose Your DBMS Product(s)! 494

Summary 495 • Key Terms 496 • Project Questions 496

ONLINE CHAPTER: SEE PAGE 495 FOR INSTRUCTIONS

Chapter 10A: Managing Databases with Microsoft SQL Server 2014

Chapter Objectives

The Microsoft SQL Server 2014 DBMS

Installing Microsoft SQL Server 2014

Microsoft SQL Server 2014 Utilities

SQL CMD and Microsoft PowerShell • *Microsoft SQL CLR* • *SQL Server 2014 Management Studio*

Creating a Microsoft SQL Server 2014 Database

Microsoft SQL Server 2014 SQL Statements and SQL Scripts

Using Existing SQL Scripts • *Using a Single SQL Script to Store Multiple SQL Commands*

Creating and Implementing the View Ridge Gallery VRG Database in Microsoft SQL Server 2014

Using SQL Scripts to Create and Populate Database Tables • *Creating the View Ridge Database Table Structure* • *Reviewing Database Structures in the SQL Server GUI Display* • *Indexes* • *Populating the VRG Database Tables with Data* • *Creating SQL Views*

Importing Microsoft Excel Data into a Microsoft SQL Server Database Table

Microsoft SQL Server 2014 Application Logic

Transact-SQL • *User-Defined Functions* • *Stored Procedures* • *Triggers*

Microsoft SQL Server 2014 Concurrency Control

Transaction Isolation Level • *Cursor Concurrency* • *Locking Hints*

Microsoft SQL Server 2014 Security

SQL Server Database Security Settings

Microsoft SQL Server 2014 Backup and Recovery

Backing Up a Database • *SQL Server Recovery Models* • *Restoring a Database* • *Database Maintenance Plans*

Topics Not Discussed in This Chapter

[Summary](#) • [Key Terms](#) • [Review Questions](#) • [Project Questions](#) • [Case Questions](#) • [The Queen Anne Curiosity Shop](#) • [Morgan Importing](#)

ONLINE CHAPTER: SEE PAGE 495 FOR INSTRUCTIONS

Chapter 10B: Managing Databases with Oracle Database

Chapter Objectives

The Oracle Corporation Oracle Database

Installing a Loopback Adapter • *Oracle Database, Java, JavaScript, and the Adobe Flash Player* • *Oracle Database Documentation*

Installing Oracle Database 12c with the Oracle Universal Installer (OUI)

Installing Oracle Database Express Edition 11g Release 2 (Oracle Database XE)

Oracle Database Database Administration Tools

The Oracle Database 12c Configuration Assistant • *The Oracle Enterprise Manager Database Express 12c Database Administration Utility* • *The Oracle Database XE 11.2 Database Administration Utility*

Oracle Tablespaces

Oracle Database Security

User Privileges • *Creating a User Account* • *Creating a Role*

Oracle Database Application Development Tools

*Oracle SQL*Plus* • *Oracle SQL Developer* • *Creating a Workspace for the SQL Developer Files* • *Oracle Database Schemas*

Creating and Using a Oracle Database Database

Creating a Database in Oracle Database • *Oracle Database SQL Statements and SQL Scripts* • *Using Existing SQL Scripts* • *Using a Single SQL Script to Store Multiple SQL Commands*

Implementing the View Ridge Gallery VRG Database in Oracle Database

Using SQL Scripts to Create and Populate Database Tables • *Creating the View Ridge Gallery VRG Database Table Structure* • *Transaction COMMIT in Oracle Database* • *Reviewing Database Structures in the SQL Developer GUI Display* • *Indexes* • *Populating the VRG Tables* • *Creating SQL Views*

Importing Microsoft Excel Data into an Oracle Database Table

Oracle Database Application Logic

Oracle Database PL/SQL • *User-Defined Functions* • *Stored Procedures* • *Triggers*

Oracle Database Concurrency Control

Read-Committed Transaction Isolation Level • *Serializable Transaction Isolation Level* • *Read-Only Transaction Isolation* • *Additional Locking Comments*

Oracle Database Backup and Recovery

Oracle Recovery Facilities • *Types of Failure*

Topics Not Discussed in This Chapter

[Summary](#) • [Key Terms](#) • [Review Questions](#) • [Project Questions](#) • [Case Questions](#) • [The Queen Anne Curiosity Shop](#) • [Morgan Importing](#)

ONLINE CHAPTER: SEE PAGE 495 FOR INSTRUCTIONS

Chapter 10C: Managing Databases with MySQL 5.6

Chapter Objectives

The Oracle MySQL 5.6 DBMS

Installing MySQL 5.6

Configuring Non-Windows Versions of MySQL Community Server • *MySQL Storage Engines*

The MySQL Utilities

The MySQL Command-Line Client • *The MySQL Workbench GUI Utility* • *Creating a Workspace for the MySQL Workbench Files*

Creating and Using a MySQL Database

Creating a Database in MySQL • *Setting the Active Database in MySQL*

MySQL SQL Statements and SQL Scripts

Using Existing SQL Scripts • *Using a Single SQL Script to Store Multiple SQL Commands*

Implementing the View Ridge Gallery VRG Database in MySQL 5.6

Using SQL Scripts to Create and Populate Database Tables • *Creating the VRG*

Database Table Structure • *Reviewing Database Structures in the MySQL GUI*

Display • *Indexes* • *Populating the VRG Tables with Data* • *Transaction COMMIT in*

MySQL • *Creating SQL Views*

Importing Microsoft Excel Data into a MySQL 5.6 Database Table

MySQL Application Logic

MySQL SQL/PSM Procedural Statements • *User-Defined Functions* • *Stored*

Procedures • *Triggers* • *A Last Word on MySQL Stored Procedures and Triggers*

Concurrency Control

MySQL 5.6 Security

MySQL Database Security Settings

MySQL 5.6 DBMS Backup and Recovery

Backing Up a MySQL Database • *Restoring a MySQL Database*

Topics Not Discussed in This Chapter

[Summary](#) • [Key Terms](#) • [Review Questions](#) • [Project Questions](#) • [Case Questions](#) • [The Queen Anne Curiosity Shop](#) • [Morgan Importing](#)

PART 5 ■ Database Access Standards

497

Chapter 11: The Web Server Environment 498

Chapter Objectives 498

A Web Database Application for the View Ridge Gallery 500

The Web Database Processing Environment 501

Database Server Access Standards 502

The ODBC Standard 503

ODBC Architecture 504 • *Conformance Levels* 505 • *Creating an ODBC Data Source Name* 506

The Microsoft .NET Framework and ADO.NET 512

OLE DB 514 • *ADO and ADO.NET* 517 • *The ADO.NET Object Model* 518

The Java Platform 522

JDBC 522 • *Java Server Pages (JSP) and Servlets* 524 • *Apache Tomcat* 525

Web Database Processing with PHP 526

Web Database Processing with PHP and the NetBeans IDE 527 • *Getting Started with HTML Web Pages* 529 • *The index.html Web Page* 530 • *Creating the index.html Web Page* 530 • *Using PHP* 532

Web Page Examples with PHP 539

Example 1: Updating a Table 540 • *Example 2: Using PHP Data Objects (PDO)* 545 • *Example 3: Invoking a Stored Procedure* 545 • *Challenges for Web Database Processing* 551 • *SQL Injection Attacks* 552

Extensible Markup Language (XML) 552

The Importance of XML 552 • *XML as a Markup Language* 553

Creating XML Documents from Database Data 554

Using the SQL SELECT...FOR XML Statement 554

[Summary](#) 556 • [Key Terms](#) 557 • [Review Questions](#) 558 • [Project Questions](#) 561 • [Case Questions](#) 563 • [The Queen Anne Curiosity Shop](#) 563 • [Morgan Importing](#) 564

Chapter 12: Big Data, Data Warehouses, and Business Intelligence Systems 565

- Chapter Objectives 565
- Business Intelligence Systems 567
- The Relationship Between Operational and BI Systems 567
- Reporting Systems and Data Mining Applications 567
 - Reporting Systems* 567 • *Data Mining Applications* 568
- Data Warehouses and Data Marts 568
 - Components of a Data Warehouse* 569 • *Data Warehouses Versus Data Marts* 571
 - *Dimensional Databases* 573
- Reporting Systems 580
 - RFM Analysis* 580 • *OLAP* 582
- Data Mining 591
- Distributed Database Processing 592
 - Types of Distributed Databases* 592 • *Challenges of Distributed Databases* 593
- Object-Relational Databases 594
- Virtualization 595
- Cloud Computing 596
- Big Data and the Not Only SQL Movement 598
 - Column Family Databases* 598 • *MapReduce* 601 • *Hadoop* 602

Summary 602 • Key Terms 603 • Review Questions 604 • Project Questions 606 • Case Questions 607 • The Queen Anne Curiosity Shop 608 • Morgan Importing 609

Appendices

ONLINE APPENDICES: SEE PAGE 610 FOR INSTRUCTIONS

Appendix A: Getting Started with Microsoft Access 2013

- Chapter Objectives
- What Is the Purpose of This Appendix?
- Why Should I Learn to Use Microsoft Access 2013?
- What Will This Appendix Teach Me?
- What Is a Table Key?
- Relationships Among Tables
- Creating a Microsoft Access Database
- The Microsoft Office Fluent User Interface
 - The Ribbon and Command Tabs* • *Contextual Command Tabs* • *Modifying the Quick Access Toolbar* • *Database Objects and the Navigation Pane*
- Closing a Database and Exiting Microsoft Access
- Opening an Existing Microsoft Access Database
- Creating Microsoft Access Database Tables
- Inserting Data into Tables—The Datasheet View
 - Modifying and Deleting Data in Tables in the Datasheet View*
- Creating Relationships Between Tables
- Working with Microsoft Access Queries
- Microsoft Access Forms and Reports
- Closing a Database and Exiting Microsoft Access 2013

Key Terms • Review Questions

Appendix B: Getting Started with Systems Analysis and Design

Chapter Objectives

What Is the Purpose of This Appendix?

What Is Information?

What Is an Information System?

What Is a Competitive Strategy?

How Does a Company Organize Itself Based on Its Competitive Strategy?

What Is a Business Process?

How Do Information Systems Support Business Processes?

Do Information Systems Include Processes?

Do We Have to Understand Business Processes in Order to Create Information Systems?

What Is Systems Analysis and Design?

What Are the Steps in the SDLC?

The System Definition Step • The Requirements Analysis Step • The Component Design Step • The Implementation Step • The System Maintenance Step

What SDLC Details Do We Need to Know?

What Is Business Process Modeling Notation?

What Is Project Scope?

How Do I Gather Data and Information About System Requirements?

How Do Use Cases Provide Data and Information About System Requirements?

The Highline University Database

The College Report • The Department Report • The Department/Major Report • The Student Acceptance Letter

What Are Business Rules?

What Is a User Requirements Document (URD)?

What Is a Statement of Work (SOW)?

Key Terms • Review Questions • Project Questions

Appendix C: E-R Diagrams and the IDEF1X Standard

Chapter Objectives

IDEF1X Entities

IDEF1X Relationships

Nonidentifying Connection Relationships • Identifying Connection Relationships • Nonspecific Relationships • Categorization Relationships

Domains

Domains Reduce Ambiguity • Domains Are Useful • Base Domains and Typed Domains

Key Terms • Review Questions

Appendix D: E-R Diagrams and the UML Standard

Chapter Objectives

UML Entities and Relationships

Representation of Weak Entities

Representation of Subtypes

OOP Constructs Introduced by UML

The Role of UML in Database Processing Today

Key Terms • Review Questions

Appendix E: Getting Started with the MySQL Workbench Data Modeling Tools

Chapter Objectives

What Is the Purpose of This Appendix?

Why Should I Learn to Use the MySQL Workbench Data Modeling Tools?

What Will This Appendix Teach Me?

What Won't This Appendix Teach Me?

How Do I Start the MySQL Workbench?

How Do I Create a Workspace for the MySQL Workbench Files?

How Do I Install the MySQL Connector/ODBC?

How Do I Create Database Designs in the MySQL Workbench?

How Do I Create a Database Model and E-R Diagram in the MySQL Workbench?

[Key Terms](#) • [Review Questions](#) • [Project Questions](#)

Appendix F: Getting Started with Microsoft Visio 2013

Chapter Objectives

What Is the Purpose of This Appendix?

Why Should I Learn to Use Microsoft Visio 2013?

What Will This Appendix Teach Me?

What Won't This Appendix Teach Me?

How Do I Start Microsoft Visio 2013?

How Do I Create a Database Model Diagram in Microsoft Visio 2013?

How Do I Name and Save a Database Model Diagram in Microsoft Visio 2013?

How Do I Create Entities/Tables in a Database Model Diagram in Microsoft Visio 2013?

How Do I Create Relationships Between Tables in a Database Model Diagram in Microsoft Visio 2013?

[Key Terms](#) • [Review Questions](#) • [Project Questions](#)

Appendix G: Data Structures for Database Processing

Chapter Objectives

What Is the Purpose of This Appendix?

What Will This Appendix Teach Me?

What Are Flat Files?

Processing Flat Files in Multiple Orders • *A Note on Record Addressing* • *How Can Linked Lists Be Used to Maintain Logical Record Order?* • *How Can Indexes Be Used to Maintain Logical Record Order?* • *B-Trees* • *Summary of Data Structures*

How Can We Represent Binary Relationships?

A Review of Record Relationships • *How Can We Represent Trees?* • *How Can We Represent Simple Networks?* • *How Can We Represent Complex Networks?* • *Summary of Relationship Representations*

How Can We Represent Secondary Keys?

How Can We Represent Secondary Keys with Linked-Lists? • *How Can We Represent Secondary Keys with Indexes?*

[Key Terms](#) • [Review Questions](#)

Appendix H: The Semantic Object Model

Chapter Objectives

What Is the Purpose of This Appendix?

Why Should I Learn to Use the Semantic Object Model?

What Will This Appendix Teach Me?

What Are Semantic Objects?

What Semantic Objects Are Used in the Semantic Object Model?

What Are Semantic Object Attributes? • *What Are Object Identifiers?* • *What Are Attribute Domains?* • *What Are Semantic Object Views?*

What Types of Objects Are Used in the Semantic Object Model?

What Are Simple Objects? • *What Are Composite Objects?* • *What Are Compound Objects?* • *How Do We Represent One-to-One Compound Objects as Relational Structures?* • *How Do We Represent One-to-Many and Many-to-One Relationships as Relational Structures?* • *How Do We Represent Many-to-Many Relationship Objects as Relational Structures?* • *What Are Hybrid Objects?* • *How Do We Represent Hybrid Object Relationships as Relational Structures?* • *What Are Association Objects?* • *What Are Parent/Subtype Objects?* • *What Are Archetype/Version Objects?*

Comparing the Semantic Object and the E-R Models

Key Terms • **Review Questions**

Appendix I: Getting Started with Web Servers, PHP, and the NetBeans IDE

Chapter Objectives

What Is the Purpose of This Appendix?

How Do I Install a Web Server?

How Do I Set Up IIS in Windows 8.1?

How Do I Manage IIS in Windows 8.1?

How Is a Web Site Structured?

How Do I View a Web Page from the IIS Web Server?

How Is Web Site Security Managed?

What Is the NetBeans IDE?

How Do I Install the NetBeans IDE?

What Is PHP?

How Do I Install PHP?

How Do I Create a Web Page Using the NetBeans IDE

How Do I Manage the PHP Configuration?

Key Terms • **Review Questions** • **Project Questions**

Appendix J: Business Intelligence Systems

Chapter Objectives

What Is the Purpose of This Appendix?

Business Intelligence Systems

Reporting Systems and Data Mining Applications

Reporting Systems • *Data Mining Applications*

The Components of a Data Warehouse

Data Warehouses and Data Marts • *Data Warehouses and Dimensional Databases*

Reporting Systems

RFM Analysis • *Producing the RFM Report* • *Reporting System Components* • *Report Types* • *Report Media* • *Report Modes* • *Report System Functions* • *OLAP*

Data Mining

Unsupervised Data Mining • *Supervised Data Mining* • *Three Popular Data Mining Techniques* • *Market Basket Analysis* • *Using SQL for Market Basket Analysis*

Key Terms • **Review Questions** • **Project Questions**

Appendix K: Big Data

Chapter Objectives

What Is the Purpose of This Appendix?

What Is Big Data?

The Three Vs and the “Wanna Vs”

Big Data and NoSQL Systems

The CAP Theorem

Extensible Markup Language (XML)

XML as a Markup Language • *XML Schema* • *Creating XML Documents from Database Data* • *Why Is XML Important?* • *Additional XML Standards*

Non-Relational Database Management Systems

Key Value Databases • *Document Databases* • *Column Family Databases* • *Graph Databases*

Big Data, NoSQL Systems, and the Future

Key Terms • **Review Questions** • **Project Questions**

Bibliography 611

Glossary 613

Index 628

The 14th Global Edition of *Database Processing: Fundamentals, Design, and Implementation* refines the organization and content of this classic textbook to reflect a new teaching and professional workplace environment. Students and other readers of this book will benefit from new content and features in this edition.

New to This Edition

Content and features new to the 14th Global Edition of *Database Processing: Fundamentals, Design, and Implementation* include the following:

- The SQL topics in Chapter 2 have been reorganized and expanded to provide a more concise presentation of SQL queries. New material on SQL set operators (UNION, INTERSECTION, and EXCEPT) has been added to ensure that nearly all SQL query topics are covered in one chapter (the exception is correlated subqueries, which are still reserved for Chapter 8).
- The material on *Big Data* and the evolving *NoSQL movement* is summarized in Chapter 12 and then expanded upon in a new Appendix K—*Big Data*. This is an important topic that is constantly developing and changing, and the new appendix provides room for an extended discussion of the topic. Material on virtualization and cloud computing is updated in Chapter 12.
- Online chapters on Microsoft SQL Server 2014 (Chapter 10A), Oracle Database (Chapter 10B), and MySQL 5.6 (Chapter 10C) now have a section on importing data from Microsoft Excel 2013 worksheets.
- The book has been updated to reflect the use of Microsoft SQL Server 2014, the current version of Microsoft SQL Server. Although most of the topics covered are backward compatible with Microsoft SQL Server 2012 and Microsoft SQL Server 2008 R2 Express edition, all material in the book now uses SQL Server 2014 in conjunction with Office 2013 exclusively.
- Oracle's Oracle Database is now updated to Oracle Database 12c, and Oracle Database Express Edition 11g Release 2 (Oracle Database XE) is introduced as the preferred Oracle Database product for use on personal computers. The current version of the Oracle SQL Developer GUI tool provides a common interface to both versions of Oracle Database, and we provide detailed examples of how to use it.
- Microsoft Windows Server 2012 R2 is the server operating system and Windows 8.1 is the workstation operating system generally discussed and illustrated in the text. These are the current Microsoft server and workstation operating systems.
- We have updated online Appendix I—*Getting Started with Web Servers, PHP and the NetBeans IDE*. We are now using the NetBeans IDE instead of the Eclipse PDT IDE. This provides a better development environment with a much simpler set of product installations because the Java JDK and NetBeans are installed in one combined installation. This new material provides a simplified (but still detailed) introduction to the installation and use of the Microsoft IIS Web server, PHP, the Java JDK, and the NetBeans in Appendix I. All of these tools are then used for Web database application development as discussed in Chapter 11.

Fundamentals, Design, and Implementation

With today's technology, it is impossible to utilize a DBMS successfully without first learning fundamental concepts. After years of developing databases with business users, we have developed what we believe to be a set of essential database concepts. These are augmented by the concepts necessitated by the increasing use of the Internet, the World Wide Web, and commonly available analysis tools. Thus, the organization and topic selection of the 14th Global Edition are designed to:

- Present an early introduction to SQL queries.
- Use a “spiral approach” to database design.
- Use a consistent, generic Information Engineering (IE) Crow's Foot E-R diagram notation for data modeling and database design.
- Provide a detailed discussion of specific normal forms within a discussion of normalization that focuses on pragmatic normalization techniques.
- Use current DBMS technology: Microsoft Access 2013, Microsoft SQL Server 2014, Oracle Database 12c (and alternately Oracle Database Express Edition 11g Release 2), and MySQL 5.6.
- Create Web database applications based on widely used Web development technology.
- Provide an introduction to business intelligence (BI) systems.
- Discuss the dimensional database concepts used in database designs for data warehouses and online analytical processing (OLAP).
- Discuss the emerging and important topics of server virtualization, cloud computing, Big Data, and the NoSQL (Not only SQL) movement.

These changes have been made because it has become obvious that the basic structure of the earlier editions (up to and including the 9th edition—the 10th edition introduced many of the changes we used in the 11th, 12th, and 13th editions and retain in the 14th edition) was designed for a teaching environment that no longer exists. The structural changes to the book were made for several reasons:

- Unlike the early years of database processing, today's students have ready access to data modeling and DBMS products.
- Today's students are too impatient to start a class with lengthy conceptual discussions on data modeling and database design. They want to do something, see a result, and obtain feedback.
- In the current economy, students need to reassure themselves that they are learning marketable skills.

Early Introduction of SQL DML

Given these changes in the classroom environment, this book provides an early introduction to SQL data manipulation language (DML) SELECT statements. The discussion of SQL data definition language (DDL) and additional DML statements occurs in Chapters 7 and 8. By encountering SQL SELECT statements in Chapter 2, students learn early in the class how to query data and obtain results, seeing firsthand some of the ways that database technology will be useful to them.

The text assumes that students will work through the SQL statements and examples with a DBMS product. This is practical today because nearly every student has access to Microsoft Access. Therefore, Chapters 1 and 2 and Appendix A—*Getting Started with Microsoft Access 2013*, are written to support an early introduction of Microsoft Access 2013 and the use of Microsoft Access 2013 for SQL queries (Microsoft Access 2013 QBE query techniques are also covered).

If a non-Microsoft Access-based approach is desired, versions of Microsoft SQL Server 2014, Oracle Database, and MySQL 5.6 are readily available for use. Free versions of the three major DBMS products covered in this book (SQL Server 2014 Express Edition, Oracle Database Express Edition 11g Release 2 (Oracle Database XE), and MySQL 5.6 Community Edition) are available for download. Thus, students can actively use a DBMS product by the end of the first week of class.

BY THE WAY

The presentation and discussion of SQL are spread over four chapters so students can learn about this important topic in small bites. SQL SELECT statements are taught in Chapter 2. SQL data definition language (DDL) and SQL data manipulation language (DML) statements are presented in Chapter 7. Correlated subqueries and EXISTS/NOT EXISTS statements are described in Chapter 8, while SQL transaction control language (TCL) and SQL data control language (DCL) are discussed in Chapter 9. Each topic appears in the context of accomplishing practical tasks. Correlated subqueries, for example, are used to verify functional dependency assumptions, a necessary task for database redesign.

This box illustrates another feature used in this book: BY THE WAY boxes are used to separate comments from the text discussion. Sometimes they present ancillary material; other times they reinforce important concepts.

A Spiral Approach to the Database Design Process

Today, databases arise from three sources: (1) from the need to integrate existing data from spreadsheets, data files, and database extracts; (2) from the need to develop new information systems projects; and (3) from the need to redesign an existing database to adapt to changing requirements. We believe that the fact that these three sources exist presents instructors with a significant pedagogical opportunity. Rather than teach database design just once from data models, why not teach database design three times, once for each of these sources? In practice, this idea has turned out to be even more successful than expected.

Database Design Iteration 1: Databases from Existing Data

Considering the design of databases from existing data, if someone were to email us a set of tables and say, “Create a database from them,” how would we proceed? We would examine the tables in light of normalization criteria and then determine whether the new database was for a production system that allows new data to be inserted for each new transaction, or for a business intelligence (BI) data warehouse that allow users to only query data for use in reports and data analysis. Depending on the answer, we would normalize the data, pulling them apart (for the production transaction processing system), or denormalize the data, joining them together (for the BI system data warehouse). All of this is important for students to know and understand.

Therefore, the first iteration of database design gives instructors a rich opportunity to teach normalization, not as a set of theoretical concepts but rather as a useful toolkit for making design decisions for databases created from existing data. Additionally, the construction of databases from existing data is an increasingly common task that is often assigned to junior staff members. Learning how to apply normalization to the design of databases from existing data not only provides an interesting way of teaching normalization, it is also common and useful!

We prefer to teach and use a pragmatic approach to normalization and present this approach in Chapter 3. However, we are aware that many instructors like to teach normalization in the context of a step-by-step normal form presentation (1NF, 2NF, 3NF, then BCNF), and Chapter 3 now includes additional material to provide more support for this approach as well.

In today’s workplace, large organizations are increasingly licensing standardized software from vendors such as SAP, Oracle, and Siebel. Such software already has a database design. But with every organization running the same software, many are learning that they can gain a competitive advantage only if they make better use of the data in those pre-designed databases. Hence, students who know how to extract data and create read-only databases for reporting and data mining have obtained marketable skills in the world of ERP and other packaged software solutions.

Database Design Iteration 2: Data Modeling and Database Design

The second source of databases is from new systems development. Although not as common as in the past, many databases are still created from scratch. Thus, students still need to

learn data modeling, and they still need to learn how to transform data models into database designs that are then implemented in a DBMS product.

The IE Crow's Foot Model as a Design Standard

This edition uses a generic, standard IE Crow's Foot notation. Your students should have no trouble understanding the symbols and using the data modeling or database design tool of your choice.

IDEF1X (which was used as the preferred E-R diagram notation in the 9th edition of this text) is explained in Appendix C—*E-R Diagrams and the IDEF1X Standard*, in case your students will graduate into an environment where it is used or if you prefer to use it in your classes. UML is explained in Appendix D—*E-R Diagrams and the UML Standard*, in case you prefer to use UML in your classes.

BY THE WAY

The choice of a data modeling tool is somewhat problematic. Of the two most readily available tools, Microsoft Visio 2013 has been rewritten as a very rudimentary database design tool, while Oracle's MySQL Workbench is a database design tool, not a data modeling tool. MySQL Workbench cannot produce an N:M relationship as such (as a data model requires) but has to immediately break it into two 1:N relationships (as database design does). Therefore, the intersection table must be constructed and modeled. This confounds data modeling with database design in just the way that we are attempting to teach students to avoid.

To be fair to Microsoft Visio 2013, it is true that data models with N:M relationships can be drawn using the standard Microsoft Visio 2013 drawing tools. Unfortunately, Microsoft has chosen to remove many of the best database design tools that were in Microsoft Visio 2010, and Microsoft Visio 2013 lacks the tools that made it a favorite of Microsoft Access and Microsoft SQL Server users. For a full discussion of these tools, see Appendix E—*Getting Started with the MySQL Workbench Data Modeling Tools*, and Appendix F—*Getting Started with Microsoft Visio 2013*.

Good data modeling tools are available, but they tend to be more complex and expensive. Two examples are Visible Systems' Visible Analyst and CA Technologies' CA ERwin Data Modeler. Visible Analyst is available in a student edition (at a modest price), and a one-year time-limited CA Technologies' ERwin Data Modeler Community Edition suitable for class use can be downloaded from <http://erwin.com/products/data-modeler/community-edition>. CA Technologies has limited the number of objects that can be created by this edition to 25 entities per model and disabled some other features (see <http://erwin.com/content/products/CA-ERwin-r9-Community-Edition-Matrix-na.pdf>), but there is still enough functionality to make this product a possible choice for class use.

Database Design from E-R Data Models

As we discuss in Chapter 6, designing a database from data models consists of three tasks: representing entities and attributes with tables and columns; representing maximum cardinality by creating and placing foreign keys; and representing minimum cardinality via constraints, triggers, and application logic.

The first two tasks are straightforward. However, designs for minimum cardinality are more difficult. Required parents are easily enforced using NOT NULL foreign keys and referential integrity constraints. Required children are more problematic. In this book, however, we simplify the discussion of this topic by limiting the use of referential integrity actions and by supplementing those actions with design documentation. See the discussion around Figure 6-29.

Although the design for required children is complicated, it is important for students to learn. It also provides a reason for students to learn about triggers as well. In any case, the discussion of these topics is much simpler than it was in prior editions because of the use of the IE Crow's Foot model and ancillary design documentation.

Database Implementation from Database Designs

Of course, to complete the process, a database design must be implemented in a DBMS product. This is discussed in Chapter 7, where we introduce SQL DDL for creating tables and SQL DML for populating the tables with data.

BY THE WAY

David Kroenke is the creator of the semantic object model (SOM). The SOM is presented in Appendix H—*The Semantic Object Model*. The E-R data model is used everywhere else in the text.

Database Design Iteration 3: Database Redesign

Database redesign, the third iteration of database design, is both common and difficult. As stated in Chapter 8, information systems cause organizational change. New information systems give users new behaviors, and as users behave in new ways, they require changes in their information systems.

Database redesign is by nature complex. Depending on your students, you may wish to skip it, and you can do so without loss of continuity. Database redesign is presented after the discussion of SQL DDL and DML in Chapter 7 because it requires the use of advanced SQL. It also provides a practical reason to teach correlated subqueries and EXISTS/NOT EXISTS statements.

Active Use of a DBMS Product

We assume that students will actively use a DBMS product. The only real question becomes “which one?” Realistically, most of us have four alternatives to consider: Microsoft Access, Microsoft SQL Server, Oracle Database, and MySQL. You can use any of those products with this text, and tutorials for each of them are presented for Microsoft Access 2013 (Appendix A), SQL Server 2014 (Chapter 10A), Oracle Database 12c and Oracle Database Express Edition 11g Release 2 (Chapter 10B), and MySQL 5.6 (Chapter 10C). Given the limitations of class time, it is probably necessary to pick and use just one of these products. You can often devote a portion of a lecture to discussing the characteristics of each, but it is usually best to limit student work to one of them. The possible exception to this is starting the course with Microsoft Access and then switching to a more robust DBMS product later in the course.

Using Microsoft Access 2013

The primary advantage of Microsoft Access is accessibility. Most students already have a copy, and, if not, copies are easily obtained. Many students will have used Microsoft Access in their introductory or other classes. Appendix A—*Getting Started with Microsoft Access 2013* is a tutorial on Microsoft Access 2013 for students who have not used it but who wish to use it with this book.

However, Microsoft Access has several disadvantages. First, as explained in Chapter 1, Microsoft Access is a combination application generator and DBMS. Microsoft Access confuses students because it confounds database processing with application development. Also, Microsoft Access 2013 hides SQL behind its query processor and makes SQL appear as an afterthought rather than a foundation. Furthermore, as discussed in Chapter 2, Microsoft Access 2013 does not correctly process some of the basic SQL-92 standard statements in its default setup. Finally, Microsoft Access 2013 does not support triggers. You can simulate triggers by trapping Windows events, but that technique is nonstandard and does not effectively communicate the nature of trigger processing.

Using Microsoft SQL Server 2014, Oracle Database, or MySQL 5.6

Choosing which of these products to use depends on your local situation. Oracle Database 12c, a superb enterprise-class DBMS product, is difficult to install and administer. However, if you have local staff to support your students, it can be an excellent choice. Fortunately, Oracle Database Express Edition 11g Release 2, commonly referred to as Oracle Database XE, is easy to install, easy to use, and freely downloadable. If you want

your students to be able to install Oracle Database on their own computers, use Oracle Database XE. As shown in Chapter 10B, Oracle's SQL Developer GUI tool (or SQL*Plus if you are dedicated to this beloved command-line tool) is a handy tool for learning SQL, triggers, and stored procedures.

Microsoft SQL Server 2014, although probably not as robust as Oracle Database, is easy to install on Windows machines, and it provides the capabilities of an enterprise-class DBMS product. The standard database administrator tool is the Microsoft SQL Server Management Studio GUI tool. As shown in Chapter 10A, SQL Server 2014 can be used to learn SQL, triggers, and stored procedures.

MySQL 5.6, discussed in Chapter 10C, is an open source DBMS product that is receiving increased attention and market share. The capabilities of MySQL are continually being upgraded, and MySQL 5.6 supports stored procedures and triggers. MySQL also has excellent GUI tools in the MySQL Workbench and an excellent command-line tool (the MySQL Command Line Client). It is the easiest of the three products for students to install on their own computers. It also works with the Linux operating system and is popular as part of the AMP (Apache-MySQL-PHP) package (known as WAMP on Windows and LAMP on Linux).

BY THE WAY

Because we only present currently available software products in this book, we cover MySQL 5.6 instead of MySQL 5.7. However, MySQL 5.7 is currently in release candidate status, which means that it will be generally available in the near future. All discussion of MySQL 5.6 in this book will also apply to MySQL 5.7.

BY THE WAY

If the DBMS you use is not driven by local circumstances and you do have a choice, we recommend using Microsoft SQL Server 2014. It has all of the features of an enterprise-class DBMS product, and it is easy to install and use. Another option is to start with Microsoft Access 2013 if it is available and switch to SQL Server 2014 at Chapter 7. Chapters 1 and 2 and Appendix A are written specifically to support this approach. A variant is to use Microsoft Access 2013 as the development tool for forms and reports running against an SQL Server 2014 database.

If you prefer a different DBMS product, you can still start with Microsoft Access 2013 and switch later in the course. See the detailed discussion of the available DBMS products in Chapter 10 for a good review of your options.

Focus on Database Application Processing

In this edition, we clearly draw the line between *application development* per se and *database application processing*. Specifically, we have:

- Focused on specific database dependent applications:
 - Web-based, database-driven applications
 - XML-based data processing
 - Business intelligence (BI) systems applications
- Emphasized the use of commonly available, multiple-OS-compatible application development languages.
- Limited the use of specialized vendor-specific tools and programming languages as much as possible.

There is simply not enough room in this book to provide even a basic introduction to programming languages used for application development such as the Microsoft .NET languages and Java. Therefore, rather than attempting to introduce these languages, we leave them for other classes where they can be covered at an appropriate depth. Instead, we focus on

basic tools that are relatively straightforward to learn and immediately applicable to database-driven applications. We use PHP as our Web development language, and we use the readily available NetBeans integrated development environment (IDE) as our development tool. The result is a very focused final section of the book, where we deal specifically with the interface between databases and the applications that use them.

BY THE WAY

Although we try to use widely available software as much as possible, there are, of course, exceptions where we must use vendor-specific tools. For BI applications, for example, we draw on Microsoft Excel's PivotTable capabilities and the Microsoft PowerPivot for Microsoft Excel 2013 add-in and on the Microsoft SQL Server 2012 R2 Data Mining Add-ins for Microsoft Office. However, either alternatives to these tools are available (OpenOffice.org DataPilot capabilities, the Palo OLAP Server) or the tools are generally available for download.

Business Intelligence Systems and Dimensional Databases

This edition maintains coverage of business intelligence (BI) systems (Chapter 12 and Appendix J). The chapter includes a discussion of dimensional databases, which are the underlying structure for data warehouses, data marts, and OLAP servers. It still covers data management for data warehouses and data marts and also describes reporting and data mining applications, including OLAP.

Appendix J includes in-depth coverage of two applications that should be particularly interesting to students. The first is RFM analysis, a reporting application frequently used by mail order and e-commerce companies. The complete RFM analysis is accomplished in Appendix J through the use of standard SQL statements. This chapter can be assigned at any point after Chapter 8 and could be used as a motivator to illustrate the practical applications of SQL mid-course. Finally, Appendix K provides additional material on Big Data and NoSQL databases to supplement and support Chapter 12.

Overview of the Chapters in the 14th Global Edition

Chapter 1 sets the stage by introducing database processing, describing basic components of database systems, and summarizing the history of database processing. If students are using Microsoft Access 2013 for the first time (or need a good review), they will also need to study Appendix A—*Getting Started with Microsoft Access 2013* at this point. Chapter 2 presents SQL SELECT statements. It also includes sections on how to submit SQL statements to Microsoft Access 2013, SQL Server 2014, Oracle Database, and MySQL 5.6.

The next four chapters, Chapters 3 through 6, present the first two iterations of database design. Chapter 3 presents the principles of normalization to Boyce-Codd normal form (BCNF). It describes the problems of multivalued dependencies and explains how to eliminate them. This foundation in normalization is applied in Chapter 4 to the design of databases from existing data.

Chapters 5 and 6 describe the design of new databases. Chapter 5 presents the E-R data model. Traditional E-R symbols are explained, but the majority of the chapter uses IE Crow's Foot notation. Chapter 5 provides a taxonomy of entity types, including strong, ID-dependent, weak but not ID-dependent, supertype/subtype, and recursive. The chapter concludes with a simple modeling example for a university database.

Chapter 6 describes the transformation of data models into database designs by converting entities and attributes to tables and columns, by representing maximum cardinality by creating and placing foreign keys, and by representing minimum cardinality via carefully designed DBMS constraints, triggers, and application code. The primary section of this chapter parallels the entity taxonomy in Chapter 5.

Chapter 7 presents SQL DDL, DML, and SQL/Persistent Stored Modules (SQL/PSM). SQL DDL is used to implement the design of an example introduced in Chapter 6. INSERT, UPDATE, MERGE, and DELETE statements are discussed, as are SQL views. Additionally, the principles of embedding SQL in program code are presented, SQL/PSM is discussed, and triggers and stored procedures are explained.

Database redesign, the third iteration of database design, is described in Chapter 8. This chapter presents SQL statements using correlated subqueries and the SQL EXIST and NOT EXISTS operators, and uses these statements in the redesign process. Reverse engineering is described, and basic redesign patterns are illustrated and discussed.

Chapters 9, 10, 10A, 10B, and 10C consider the management of multiuser organizational databases. Chapter 9 describes database administration tasks, including concurrency, security, and backup and recovery. Chapter 10 is a general introduction to the online Chapters 10A, 10B, and 10C, which describe SQL Server 2014, Oracle Database (both Oracle Database 12c and Oracle Database XE), and MySQL 5.6, respectively. These chapters show how to use these specific products to create database structures and process SQL statements. They also explain concurrency, security, and backup and recovery with each product. The discussion in Chapters 10A, 10B, and 10C parallels the order of discussion in Chapter 9 as much as possible, though rearrangements of some topics are made, as needed, to support the discussion of a specific DBMS product.

BY THE WAY

We have maintained or extended our coverage of Microsoft Access, Microsoft SQL Server, Oracle Database, and MySQL (introduced in *Database Processing: Fundamentals, Design, and Implementation*, 11th edition) in this book. In order to keep the bound book to a reasonable length and to keep the cost of the book down, we have chosen to provide some material by download from our Web site at www.pearsonglobaleditions.com/kroenke. There you will find:

- Chapter 10A—*Managing Databases with Microsoft SQL Server 2014*
- Chapter 10B—*Managing Databases with Oracle Database*
- Chapter 10C—*Managing Databases with MySQL 5.6*
- Appendix A—*Getting Started with Microsoft Access 2013*
- Appendix B—*Getting Started with Systems Analysis and Design*
- Appendix C—*E-R Diagrams and the IDEF1X Standard*
- Appendix D—*E-R Diagrams and the UML Standard*
- Appendix E—*Getting Started with MySQL Workbench Data Modeling Tools*
- Appendix F—*Getting Started with Microsoft Visio 2013*
- Appendix G—*Data Structures for Database Processing*
- Appendix H—*The Semantic Object Model*
- Appendix I—*Getting Started with Web Servers, PHP, and the NetBeans IDE*
- Appendix J—*Business Intelligence Systems*
- Appendix K—*Big Data*

Chapters 11 and 12 address standards for accessing databases. Chapter 11 presents ODBC, OLE DB, ADO.NET, ASP.NET, JDBC, and JavaServer Pages (JSP). It then introduces PHP (and the NetBeans IDE) and illustrates the use of PHP for the publication of databases via Web pages. This is followed by a description of the integration of XML and database technology. The chapter begins with a primer on XML and then shows how to use the FOR XML SQL statement in SQL Server.

Chapter 12 concludes the text with a discussion of BI systems, dimensional data models, data warehouses, data marts, server virtualization, cloud computing, Big Data, structured storage, and the Not only SQL movement.

Supplements

This text is accompanied by a wide variety of supplements. Please visit the text's Web site at www.pearsonglobaleditions.com/kroenke to access the instructor and student supplements described below. Please contact your Pearson sales representative for more details. All supplements were written by David Auer, Scott Vandenberg, Bob Yoder, and Darren Lim.

For Students

Many of the sample databases used in this text are available online in Microsoft Access, Microsoft SQL Server 2014, Oracle Database, and MySQL 5.6 formats.

For Instructors

At the Instructor Resource Center, www.pearsonglobaleditions.com/Kroenke, instructors can access a variety of print, digital, and presentation resources available with this text in downloadable format. Registration is simple and gives instructors immediate access to new titles and new editions. As a registered faculty member, you can download resource files and receive immediate access to and instructions for installing course management content on your campus server. In case you ever need assistance, our dedicated technical support team is ready to help with the media supplements that accompany this text. Visit <http://247pearsoned.com> for answers to frequently asked questions and toll-free user support phone numbers.

The following supplements are available for download to adopting instructors:

- Instructor's Resource Manual
- Test Bank
- TestGen[®] Computerized Test Bank
- PowerPoint Presentations

Acknowledgments

We are grateful for the support of many people in the development of this 14th edition and previous editions. Thanks to Rick Mathieu at James Madison University for interesting and insightful discussions on the database course. Professor Doug MacLachlan from the Marketing Department at the University of Washington was most helpful in understanding the goals, objectives, and technology of data mining, particularly as it pertains to marketing. Don Nilson, formerly of the Microsoft Corporation, helped us understand the importance of XML to database processing. Craig Pencil of Western Washington University helped us refine the use of the book in the classroom. Recently David Auer and Xiaofeng Chen team-taught a database class together at Western Washington University, and our interaction and discussions with Professor Chen resulted in several modifications and improvements in this book. Professor Chen also graciously allowed us to adopt some of his classroom examples for use in the books. Thanks are also due to Harold Wise of East Carolina University, Barry Flachsbart of Missouri University of Science and Technology, and Don Malzahn of Harper College for their comments and SQL code checking. Finally, thanks to Donna Auer for giving us permission to use her painting *Lake Samish (Looking Into Water)* as the cover art for this book.

In addition, we wish to thank the reviewers of this edition:

- Ann Aksut**, *Central Piedmont Community College*
- Allen Badgett**, *Oklahoma City University*
- Rich Beck**, *Washington University*
- Jeffrey J. Blessing**, *Milwaukee School of Engineering*
- Alan Brandyberry**, *Kent State University*
- Larry Booth**, *Clayton State University*
- Jason Deane**, *Virginia Polytechnic Institute and State University*
- Barry Flachsbart**, *Missouri University of Science and Technology*

Andy Green, *Kennesaw State University*
Dianne Hall, *Auburn University*
Jeff Hassett, *University of Utah*
Barbara Hewitt, *Texas A&M, Kingsville*
William Hochstettler, *Franklin University*
Margaret Hvatum, *St. Louis Community College*
Nitin Kale, *University of Southern California, Los Angeles*
Darrel Karbginsky, *Chemeketa Community College*
Johnny Li, *South University*
Lin Lin, *New Jersey Institute of Technology*
Mike Morris, *Southeastern Oklahoma State University*
Jane Perschbach, *Texas A&M University—Central Texas*
Catherine Ricardo, *Iona College*
Kevin Roberts, *DeVry University*
Ioulia Rytikova, *George Mason University*
Christelle Scharff, *Pace University*
Julian M. Scher, *New Jersey Institute of Technology*
Namchul Shin, *Pace University*
K. David Smith, *Cameron University*
M. Jane Stafford, *Columbia College—Jefferson City*
Marcia Williams, *Bellevue Community College*
Timothy Woodcock, *Texas A&M University—Central Texas*

Finally, we would like to thank Nicole Sam, our editor; Denise Weiss, our Program Manager Ilene Kahn, our Project Manager; and Sue Nodine, our Production Project Manager; for their professionalism, insight, support, and assistance in the development of this project. We would also like to thank Scott Vandenberg and Bod Yoder (both of Siena College) for their detailed comments on the final manuscript—this book would not be what it is without their extensive input. Finally, David Kroenke would like to thank his wife, Lynda, and David Auer would like to thank his wife, Donna, for their love, encouragement, and patience while this project was being completed.

David Kroenke
Seattle, Washington

David Auer
Bellingham, Washington

Pearson would like to thank the following people for their work on the Global Edition:

Contributors

Ketil Danielsen, *Molde University College*
Mathy Paesen, *ACE-GROEPT*
Nahed Azab, *The American University in Cairo*

Reviewers

Cindy Baker, *American University of Sharja*
Neerja Sethi, *Nanyang Technological University*
Franck Laurence, *Conservatoire National des Arts et Métier*