

Coronel / Morris

# Database Systems

Design, Implementation,  
and Management 11e



# DATABASE SYSTEMS

DESIGN, IMPLEMENTATION, AND MANAGEMENT

CARLOS CORONEL • STEVEN MORRIS



Australia • Brazil • Mexico • Singapore • United Kingdom • United States

This is an electronic version of the print textbook. Due to electronic rights restrictions, some third party content may be suppressed. Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. The publisher reserves the right to remove content from this title at any time if subsequent rights restrictions require it. For valuable information on pricing, previous editions, changes to current editions, and alternate formats, please visit [www.cengage.com/highered](http://www.cengage.com/highered) to search by ISBN#, author, title, or keyword for materials in your areas of interest.

**Database Systems: Design, Implementation,  
and Management, Eleventh Edition**

Carlos Coronel and Steven Morris

Product Director: Joe Sabatino

Product Manager: Clara Goosman

Senior Content Developer: Kate Mason

Content Developer: Anne Merrill

Product Assistant: Brad Sullender

Marketing Manager: Erik La Scola

Marketing Coordinator: Elizabeth Murphy

Design and Production Service: GEX Publishing  
Services

Manufacturing Planner: Ron Montgomery

Rights Acquisition Specialist: Christine Myaskovsky

Cover image(s): ©iStockphoto.com/derrrek

© 2015, 2013 Cengage Learning

WCN: 02-200-203

ALL RIGHTS RESERVED. No part of this work covered by the copyright herein may be reproduced, transmitted, stored, or used in any form or by any means graphic, electronic, or mechanical, including but not limited to photocopying, recording, scanning, digitizing, taping, Web distribution, information networks, or information storage and retrieval systems, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the publisher.

For product information and technology assistance, contact us at  
**Cengage Learning Customer & Sales Support, 1-800-354-9706**

For permission to use material from this text or product, submit all  
requests online at [www.cengage.com/permissions](http://www.cengage.com/permissions).  
Further permissions questions can be e-mailed to  
[permissionrequest@cengage.com](mailto:permissionrequest@cengage.com)

Library of Congress Control Number: 2013954240

Student Edition:

ISBN-13: 978-1-285-19614-5

ISBN-10: 1-285-19614-7

Instructor Edition:

ISBN-13: 978-1-285-19618-3

ISBN-10: 1-285-19618-X

**Cengage Learning**

200 First Stamford Place, 4th Floor  
Stamford, CT 06902  
USA

Cengage Learning is a leading provider of customized learning solutions with office locations around the globe, including Singapore, the United Kingdom, Australia, Mexico, Brazil, and Japan. Locate your local office at:

**[www.cengage.com/global](http://www.cengage.com/global)**

Cengage Learning products are represented in Canada by Nelson Education, Ltd.

For your course and learning solutions, visit

**[www.cengage.com](http://www.cengage.com)**.

Purchase any of our products at your local college store or at our preferred online store

**[www.cengagebrain.com](http://www.cengagebrain.com)**.**Instructors:** Please visit [login.cengage.com](http://login.cengage.com) and log in to access instructor-specific resources.

Screenshots for this book were created using Microsoft Access® and Visio®, and were used with permission from Microsoft. Microsoft and the Office logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Oracle is a registered trademark, and Oracle® and MySQL are trademarks of Oracle Corporation.

iPhone, iPad, and iPod are registered trademarks of Apple Inc.

**Notice to the Reader**

Publisher does not warrant or guarantee any of the products described herein or perform any independent analysis in connection with any of the product information contained herein. Publisher does not assume, and expressly disclaims, any obligation to obtain and include information other than that provided to it by the manufacturer. The reader is expressly warned to consider and adopt all safety precautions that might be indicated by the activities described herein and to avoid all potential hazards. By following the instructions contained herein, the reader willingly assumes all risks in connection with such instructions. The publisher makes no representations or warranties of any kind, including but not limited to, the warranties of fitness for particular purpose or merchantability, nor are any such representations implied with respect to the material set forth herein, and the publisher takes no responsibility with respect to such material. The publisher shall not be liable for any special, consequential, or exemplary damages resulting, in whole or part, from the readers' use of, or reliance upon, this material.

Printed in the United States of America  
1 2 3 4 5 6 7 17 16 15 14 13



**D**edication

To the treasures in my life: To Victoria, for 23 wonderful years. Thank you for your unending support, for being my angel, my sweetie, and most importantly, my best friend. To Carlos Anthony for the many good times you have given us; you show us the way. Thank you for your words of wisdom, contagious happiness, and for bringing us shining days. You are still young; your best times are still to come. To Gabriela Victoria, who is the image of brilliance, beauty, and faithfulness. Thank you for being the sunshine in my cloudy days. To Christian Javier, whose endless energy and delightful smiles keep us going every day. Thank you for being the youthful reminder of life's simple beauties. To my parents, Sarah and Carlos, thank you for your sacrifice and example. To all of you, you are all my inspiration. "TQTATA."

**Carlos Coronel**

To Pamela, from high school sweetheart through 24 years of marriage, the beautiful love of my life who has supported, encouraged, and inspired me. More than anyone else, you are responsible for whatever successes I have achieved. To my son, Alexander Logan, whose depth of character is without measure. You are my pride and joy. To my daughter, Lauren Elizabeth, whose beauty and intensity take my breath away. You are my heart and soul. Thank you all for the sacrifices you have made that enabled me to pursue this dream. I love you so much more than I can express. To my mother, Florence Maryann, and to the memory of my father, Alton Lamar, who together instilled in me the desire to learn and the passion to achieve. To my mother-in-law, Connie Duke, and to the memory of my father-in-law, Wayne Duke, who taught me to find joy in all things. To all of you, with all my love, I dedicate this book.

**Steven Morris****FOR PETER**

To longtime colleague and friend, Peter Rob: Your drive and dedication to your students started this book. Your depth of knowledge, attention to detail, and pursuit of excellence made it succeed. Your patience and guidance continue to light our path. It is our sincere hope that as we move forward, we can continue to live up to your standard. Enjoy your retirement, my friend; you have surely earned it.

Carlos Coronel and Steven Morris

# BRIEF CONTENTS

<b>PART I:</b>	<b>DATABASE CONCEPTS</b>	
	Chapter 1: Database Systems	4
	Chapter 2: Data Models	34
<b>PART II:</b>	<b>DESIGN CONCEPTS</b>	
	Chapter 3: The Relational Database Model	70
	Chapter 4: Entity Relationship (ER) Modeling	114
	Chapter 5: Advanced Data Modeling	161
	Chapter 6: Normalization of Database Tables	190
<b>PART III:</b>	<b>ADVANCED DESIGN AND IMPLEMENTATION</b>	
	Chapter 7: Introduction to Structured Query Language (SQL)	236
	Chapter 8: Advanced SQL	322
	Chapter 9: Database Design	412
<b>PART IV:</b>	<b>ADVANCED DATABASE CONCEPTS</b>	
	Chapter 10: Transaction Management and Concurrency Control	454
	Chapter 11: Database Performance Tuning and Query Optimization	486
	Chapter 12: Distributed Database Management Systems	522
	Chapter 13: Business Intelligence and Data Warehouses	559
<b>PART V:</b>	<b>DATABASES AND THE INTERNET</b>	
	Chapter 14: Database Connectivity and Web Technologies	626
<b>PART VI:</b>	<b>DATABASE ADMINISTRATION</b>	
	Chapter 15: Database Administration and Security	670
	<b>GLOSSARY</b>	715
	<b>INDEX</b>	735

The following appendixes are included online at [cengagebrain.com](http://cengagebrain.com)

<b>APPENDIX A:</b>	<b><u>DESIGNING DATABASES WITH VISIO PROFESSIONAL: A TUTORIAL</u></b>
<b>APPENDIX B:</b>	<b><u>THE UNIVERSITY LAB: CONCEPTUAL DESIGN</u></b>
<b>APPENDIX C:</b>	<b><u>THE UNIVERSITY LAB: CONCEPTUAL DESIGN VERIFICATION, LOGICAL DESIGN, AND IMPLEMENTATION</u></b>
<b>APPENDIX D:</b>	<b><u>CONVERTING AN ER MODEL INTO A DATABASE STRUCTURE</u></b>
<b>APPENDIX E:</b>	<b><u>COMPARISON OF ER MODEL NOTATIONS</u></b>
<b>APPENDIX F:</b>	<b><u>CLIENT/SERVER SYSTEMS</u></b>
<b>APPENDIX G:</b>	<b><u>OBJECT-ORIENTED DATABASES</u></b>
<b>APPENDIX H:</b>	<b><u>UNIFIED MODELING LANGUAGE (UML)</u></b>
<b>APPENDIX I:</b>	<b><u>DATABASES IN ELECTRONIC COMMERCE</u></b>
<b>APPENDIX J:</b>	<b><u>WEB DATABASE DEVELOPMENT WITH COLD FUSION</u></b>
<b>APPENDIX K:</b>	<b><u>THE HIERARCHICAL DATABASE MODEL</u></b>
<b>APPENDIX L:</b>	<b><u>THE NETWORK DATABASE MODEL</u></b>
<b>APPENDIX M:</b>	<b><u>MICROSOFT® ACCESS® TUTORIAL</u></b>
<b>APPENDIX N:</b>	<b><u>CREATING A NEW DATABASE USING ORACLE 12C</u></b>
<b>APPENDIX O:</b>	<b><u>DATA WAREHOUSE IMPLEMENTATION FACTORS</u></b>

The CourseMate Site for this text is available at [www.cengagebrain.com](http://www.cengagebrain.com).

# TABLE OF CONTENTS

## PART I DATABASE CONCEPTS

<b>Business Vignette</b>	<b>3</b>
<b>CHAPTER 1 DATABASE SYSTEMS</b>	<b>4</b>
<b>1.1 Why Databases?</b>	<b>5</b>
<b>1.2 Data vs. Information</b>	<b>5</b>
<b>1.3 Introducing the Database</b>	<b>7</b>
1.3.1 Role and Advantages of the DBMS	7
1.3.2 Types of Databases	9
<b>1.4 Why Database Design is Important</b>	<b>11</b>
<b>1.5 Evolution of File System Data Processing</b>	<b>14</b>
1.5.1 Manual File Systems	14
1.5.2 Computerized File Systems	14
1.5.3 File System Redux: Modern End-User Productivity Tools	17
<b>1.6 Problems with File System Data Processing</b>	<b>18</b>
1.6.1 Structural and Data Dependence	19
1.6.2 Data Redundancy	19
1.6.3 Lack of Design and Data-Modeling Skills	21
<b>1.7 Database Systems</b>	<b>21</b>
1.7.1 The Database System Environment	22
1.7.2 DBMS Functions	24
1.7.3 Managing the Database System: A Shift in Focus	27
<b>1.8 Preparing for Your Database Professional Career</b>	<b>28</b>
<b>Summary</b>	<b>30</b>
<b>Key Terms</b>	<b>30</b>
<b>Review Questions</b>	<b>31</b>
<b>Problems</b>	<b>32</b>
<b>CHAPTER 2 DATA MODELS</b>	<b>34</b>
<b>2.1 Data Modeling and Data Models</b>	<b>35</b>
<b>2.2 The Importance of Data Models</b>	<b>35</b>
<b>2.3 Data Model Basic Building Blocks</b>	<b>36</b>
<b>2.4 Business Rules</b>	<b>37</b>
2.4.1 Discovering Business Rules	38
2.4.2 Translating Business Rules into Data Model Components	38
2.4.3 Naming Conventions	39
<b>2.5 The Evolution of Data Models</b>	<b>39</b>
2.5.1 Hierarchical and Network Models	40
2.5.2 The Relational Model	41
2.5.3 The Entity Relationship Model	43
2.5.4 The Object-Oriented (OO) Model	46
2.5.5 Object/Relational and XML	48
2.5.6 Emerging Data Models: Big Data and NoSQL	48
2.5.7 Data Models: A Summary	53
<b>2.6 Degrees of Data Abstraction</b>	<b>56</b>
2.6.1 The External Model	57
2.6.2 The Conceptual Model	59
2.6.3 The Internal Model	59
2.6.4 The Physical Model	60



## TABLE OF CONTENTS

<b>Summary</b>	<b>62</b>
<b>Key Terms</b>	<b>62</b>
<b>Review Questions</b>	<b>63</b>
<b>Problems</b>	<b>64</b>

### PART II DESIGN CONCEPTS

<b>Business Vignette</b>	<b>69</b>
<b>CHAPTER 3 THE RELATIONAL DATABASE MODEL</b>	<b>70</b>
<b>3.1 A Logical View of Data</b>	<b>71</b>
3.1.1 Tables and Their Characteristics	71
<b>3.2 Keys</b>	<b>74</b>
3.2.1 Dependencies	74
3.2.2 Types of Keys	75
<b>3.3 Integrity Rules</b>	<b>77</b>
<b>3.4 Relational Algebra</b>	<b>79</b>
3.4.1 Formal Definitions and Terminology	79
3.4.2 Relational Set Operators	80
<b>3.5 The Data Dictionary and the System Catalog</b>	<b>88</b>
<b>3.6 Relationships within the Relational Database</b>	<b>90</b>
3.6.1 The 1:M Relationship	90
3.6.2 The 1:1 Relationship	92
3.6.3 The M:N Relationship	94
<b>3.7 Data Redundancy Revisited</b>	<b>98</b>
<b>3.8 Indexes</b>	<b>101</b>
<b>3.9 Codd's Relational Database Rules</b>	<b>102</b>
<b>Summary</b>	<b>104</b>
<b>Key Terms</b>	<b>104</b>
<b>Review Questions</b>	<b>105</b>
<b>Problems</b>	<b>108</b>
<b>CHAPTER 4 ENTITY RELATIONSHIP (ER) MODELING</b>	<b>114</b>
<b>4.1 The Entity Relationship Model (ERM)</b>	<b>115</b>
4.1.1 Entities	115
4.1.2 Attributes	116
4.1.3 Relationships	121
4.1.4 Connectivity and Cardinality	122
4.1.5 Existence Dependence	123
4.1.6 Relationship Strength	123
4.1.7 Weak Entities	126
4.1.8 Relationship Participation	128
4.1.9 Relationship Degree	131
4.1.10 Recursive Relationships	133
4.1.11 Associative (Composite) Entities	135
<b>4.2 Developing an ER Diagram</b>	<b>137</b>
<b>4.3 Database Design Challenges: Conflicting Goals</b>	<b>145</b>
<b>Summary</b>	<b>149</b>
<b>Key Terms</b>	<b>149</b>
<b>Review Questions</b>	<b>150</b>

# TABLE OF CONTENTS

<b>Problems</b>	<b>152</b>
<b>Cases</b>	<b>155</b>
<b>CHAPTER 5 ADVANCED DATA MODELING</b>	<b>161</b>
<b>5.1 The Extended Entity Relationship Model</b>	<b>162</b>
5.1.1 Entity Supertypes and Subtypes	162
5.1.2 Specialization Hierarchy	163
5.1.3 Inheritance	164
5.1.4 Subtype Discriminator	165
5.1.5 Disjoint and Overlapping Constraints	165
5.1.6 Completeness Constraint	167
5.1.7 Specialization and Generalization	168
<b>5.2 Entity Clustering</b>	<b>168</b>
<b>5.3 Entity Integrity: Selecting Primary Keys</b>	<b>170</b>
5.3.1 Natural Keys and Primary Keys	170
5.3.2 Primary Key Guidelines	170
5.3.3 When to Use Composite Primary Keys	171
5.3.4 When to Use Surrogate Primary Keys	172
<b>5.4 Design Cases: Learning Flexible Database Design</b>	<b>174</b>
5.4.1 Design Case 1: Implementing 1:1 Relationships	174
5.4.2 Design Case 2: Maintaining History of Time-Variant Data	175
5.4.3 Design Case 3: Fan Traps	178
5.4.4 Design Case 4: Redundant Relationships	180
<b>Summary</b>	<b>181</b>
<b>Key Terms</b>	<b>181</b>
<b>Review Questions</b>	<b>182</b>
<b>Problems</b>	<b>183</b>
<b>Cases</b>	<b>184</b>
<b>CHAPTER 6 NORMALIZATION OF DATABASE TABLES</b>	<b>190</b>
<b>6.1 Database Tables and Normalization</b>	<b>191</b>
<b>6.2 The Need for Normalization</b>	<b>191</b>
<b>6.3 The Normalization Process</b>	<b>195</b>
6.3.1 Conversion to First Normal Form	196
6.3.2 Conversion to Second Normal Form	200
6.3.3 Conversion to Third Normal Form	201
<b>6.4 Improving the Design</b>	<b>203</b>
<b>6.5 Surrogate Key Considerations</b>	<b>207</b>
<b>6.6 Higher-Level Normal Forms</b>	<b>208</b>
6.6.1 The Boyce-Codd Normal Form	208
6.6.2 Fourth Normal Form (4NF)	211
<b>6.7 Normalization and Database Design</b>	<b>213</b>
<b>6.8 Denormalization</b>	<b>217</b>
<b>6.9 Data-Modeling Checklist</b>	<b>220</b>
<b>Summary</b>	<b>222</b>
<b>Key Terms</b>	<b>222</b>
<b>Review Questions</b>	<b>223</b>
<b>Problems</b>	<b>224</b>

PART III ADVANCED DESIGN AND IMPLEMENTATION

<b>Business Vignette</b>	<b>235</b>
<b>CHAPTER 7 INTRODUCTION TO STRUCTURED QUERY LANGUAGE (SQL)</b>	<b>236</b>
<b>7.1 Introduction to SQL</b>	<b>237</b>
<b>7.2 Data Definition Commands</b>	<b>239</b>
7.2.1 The Database Model	239
7.2.2 Creating the Database	241
7.2.3 The Database Schema	241
7.2.4 Data Types	242
7.2.5 Creating Table Structures	245
7.2.6 SQL Constraints	249
7.2.7 SQL Indexes	252
<b>7.3 Data Manipulation Commands</b>	<b>254</b>
7.3.1 Adding Table Rows	254
7.3.2 Saving Table Changes	255
7.3.3 Listing Table Rows	256
7.3.4 Updating Table Rows	257
7.3.5 Restoring Table Contents	258
7.3.6 Deleting Table Rows	258
7.3.7 Inserting Table Rows with a Select Subquery	259
<b>7.4 Select Queries</b>	<b>260</b>
7.4.1 Selecting Rows with Conditional Restrictions	260
7.4.2 Arithmetic Operators: The Rule of Precedence	265
7.4.3 Logical Operators: AND, OR, and NOT	265
7.4.4 Special Operators	267
<b>7.5 Additional Data Definition Commands</b>	<b>271</b>
7.5.1 Changing a Column's Data Type	271
7.5.2 Changing a Column's Data Characteristics	272
7.5.3 Adding a Column	272
7.5.4 Dropping a Column	273
7.5.5 Advanced Data Updates	273
7.5.6 Copying Parts of Tables	275
7.5.7 Adding Primary and Foreign Key Designations	276
7.5.8 Deleting a Table from the Database	277
<b>7.6 Additional Select Query Keywords</b>	<b>277</b>
7.6.1 Ordering a Listing	277
7.6.2 Listing Unique Values	279
7.6.3 Aggregate Functions	279
7.6.4 Grouping Data	284
<b>7.7 Joining Database Tables</b>	<b>287</b>
7.7.1 Joining Tables with an Alias	290
7.7.2 Recursive Joins	290
<b>Summary</b>	<b>292</b>
<b>Key Terms</b>	<b>293</b>
<b>Review Questions</b>	<b>293</b>
<b>Problems</b>	<b>294</b>
<b>Cases</b>	<b>311</b>

# TABLE OF CONTENTS

<b>CHAPTER 8</b>	<b>ADVANCED SQL</b>	<b>322</b>
<b>8.1</b>	<b>SQL Join Operators</b>	<b>323</b>
8.1.1	Cross Join	324
8.1.2	Natural Join	325
8.1.3	JOIN USING Clause	327
8.1.4	JOIN ON Clause	328
8.1.5	Outer Joins	329
<b>8.2</b>	<b>Subqueries and Correlated Queries</b>	<b>332</b>
8.2.1	WHERE Subqueries	334
8.2.2	IN Subqueries	336
8.2.3	HAVING Subqueries	336
8.2.4	Multirow Subquery Operators: ANY and ALL	337
8.2.5	FROM Subqueries	338
8.2.6	Attribute List Subqueries	339
8.2.7	Correlated Subqueries	341
<b>8.3</b>	<b>SQL Functions</b>	<b>345</b>
8.3.1	Date and Time Functions	345
8.3.2	Numeric Functions	349
8.3.3	String Functions	349
8.3.4	Conversion Functions	351
<b>8.4</b>	<b>Relational Set Operators</b>	<b>354</b>
8.4.1	UNION	354
8.4.2	UNION ALL	356
8.4.3	INTERSECT	357
8.4.4	EXCEPT (MINUS)	358
8.4.5	Syntax Alternatives	359
<b>8.5</b>	<b>Virtual Tables: Creating a View</b>	<b>361</b>
8.5.1	Updatable Views	362
<b>8.6</b>	<b>Oracle Sequences</b>	<b>365</b>
<b>8.7</b>	<b>Procedural SQL</b>	<b>369</b>
8.7.1	Triggers	374
8.7.2	Stored Procedures	385
8.7.3	PL/SQL Processing with Cursors	391
8.7.4	PL/SQL Stored Functions	394
<b>8.8</b>	<b>Embedded SQL</b>	<b>394</b>
	<b>Summary</b>	<b>399</b>
	<b>Key Terms</b>	<b>400</b>
	<b>Review Questions</b>	<b>400</b>
	<b>Problems</b>	<b>401</b>
	<b>Cases</b>	<b>410</b>
<b>CHAPTER 9</b>	<b>DATABASE DESIGN</b>	<b>412</b>
<b>9.1</b>	<b>The Information System</b>	<b>413</b>
<b>9.2</b>	<b>The Systems Development Life Cycle</b>	<b>414</b>
9.2.1	Planning	415
9.2.2	Analysis	416
9.2.3	Detailed Systems Design	416
9.2.4	Implementation	417
9.2.5	Maintenance	417

## TABLE OF CONTENTS

<b>9.3 The Database Life Cycle</b>	<b>418</b>
9.3.1 The Database Initial Study	418
9.3.2 Database Design	422
9.3.3 Implementation and Loading	423
9.3.4 Testing and Evaluation	426
9.3.5 Operation	429
9.3.6 Maintenance and Evolution	429
<b>9.4 Conceptual Design</b>	<b>430</b>
9.4.1 Data Analysis and Requirements	431
9.4.2 Entity Relationship Modeling and Normalization	433
9.4.3 Data Model Verification	436
9.4.4 Distributed Database Design	439
<b>9.5 DBMS Software Selection</b>	<b>439</b>
<b>9.6 Logical Design</b>	<b>440</b>
9.6.1 Map the Conceptual Model to the Logical Model	440
9.6.2 Validate the Logical Model Using Normalization	442
9.6.3 Validate Logical Model Integrity Constraints	442
9.6.4 Validate the Logical Model Against User Requirements	443
<b>9.7 Physical Design</b>	<b>443</b>
9.7.1 Define Data Storage Organization	443
9.7.2 Define Integrity and Security Measures	444
9.7.3 Determine Performance Measures	444
<b>9.8 Database Design Strategies</b>	<b>445</b>
<b>9.9 Centralized vs. Decentralized Design</b>	<b>446</b>
<b>Summary</b>	<b>449</b>
<b>Key Terms</b>	<b>449</b>
<b>Review Questions</b>	<b>449</b>
<b>Problems</b>	<b>450</b>

## PART IV ADVANCED DATABASE CONCEPTS

<b>Business Vignette</b>	<b>453</b>
<b>CHAPTER 10 TRANSACTION MANAGEMENT AND CONCURRENCY CONTROL</b>	<b>454</b>
<b>10.1 What Is a Transaction?</b>	<b>455</b>
10.1.1 Evaluating Transaction Results	456
10.1.2 Transaction Properties	459
10.1.3 Transaction Management with SQL	460
10.1.4 The Transaction Log	460
<b>10.2 Concurrency Control</b>	<b>461</b>
10.2.1 Lost Updates	462
10.2.2 Uncommitted Data	463
10.2.3 Inconsistent Retrievals	464
10.2.4 The Scheduler	465
<b>10.3 Concurrency Control with Locking Methods</b>	<b>466</b>
10.3.1 Lock Granularity	467
10.3.2 Lock Types	470
10.3.3 Two-Phase Locking to Ensure Serializability	471
10.3.4 Deadlocks	472



# TABLE OF CONTENTS

<b>10.4 Concurrency Control with Time Stamping Methods</b>	<b>473</b>
10.4.1 Wait/Die and Wound/Wait Schemes	474
<b>10.5 Concurrency Control with Optimistic Methods</b>	<b>475</b>
<b>10.6 ANSI Levels of Transaction Isolation</b>	<b>475</b>
<b>10.7 Database Recovery Management</b>	<b>477</b>
10.7.1 Transaction Recovery	477
<b>Summary</b>	<b>481</b>
<b>Key Terms</b>	<b>482</b>
<b>Review Questions</b>	<b>482</b>
<b>Problems</b>	<b>483</b>
<b>CHAPTER 11 DATABASE PERFORMANCE TUNING AND QUERY OPTIMIZATION</b>	<b>486</b>
<b>11.1 Database Performance-Tuning Concepts</b>	<b>487</b>
11.1.1 Performance Tuning: Client and Server	488
11.1.2 DBMS Architecture	489
11.1.3 Database Query Optimization Modes	490
11.1.4 Database Statistics	491
<b>11.2 Query Processing</b>	<b>493</b>
11.2.1 SQL Parsing Phase	494
11.2.2 SQL Execution Phase	495
11.2.3 SQL Fetching Phase	495
11.2.4 Query Processing Bottlenecks	495
<b>11.3 Indexes and Query Optimization</b>	<b>496</b>
<b>11.4 Optimizer Choices</b>	<b>498</b>
11.4.1 Using Hints to Affect Optimizer Choices	500
<b>11.5 SQL Performance Tuning</b>	<b>501</b>
11.5.1 Index Selectivity	501
11.5.2 Conditional Expressions	502
<b>11.6 Query Formulation</b>	<b>504</b>
<b>11.7 DBMS Performance Tuning</b>	<b>505</b>
<b>11.8 Query Optimization Example</b>	<b>507</b>
<b>Summary</b>	<b>516</b>
<b>Key Terms</b>	<b>517</b>
<b>Review Questions</b>	<b>517</b>
<b>Problems</b>	<b>518</b>
<b>CHAPTER 12 DISTRIBUTED DATABASE MANAGEMENT SYSTEMS</b>	<b>522</b>
<b>12.1 The Evolution of Distributed Database Management Systems</b>	<b>523</b>
<b>12.2 DDBMS Advantages and Disadvantages</b>	<b>525</b>
<b>12.3 Distributed Processing and Distributed Databases</b>	<b>526</b>
<b>12.4 Characteristics of Distributed Database Management Systems</b>	<b>527</b>
<b>12.5 DDBMS Components</b>	<b>529</b>
<b>12.6 Levels of Data and Process Distribution</b>	<b>530</b>
12.6.1 Single-Site Processing, Single-Site Data	530
12.6.2 Multiple-Site Processing, Single-Site Data	531
12.6.3 Multiple-Site Processing, Multiple-Site Data	532
<b>12.7 Distributed Database Transparency Features</b>	<b>533</b>
<b>12.8 Distribution Transparency</b>	<b>534</b>
<b>12.9 Transaction Transparency</b>	<b>536</b>
12.9.1 Distributed Requests and Distributed Transactions	536

## TABLE OF CONTENTS

12.9.2	Distributed Concurrency Control	540
12.9.3	Two-Phase Commit Protocol	540
<b>12.10</b>	<b>Performance and Failure Transparency</b>	<b>541</b>
<b>12.11</b>	<b>Distributed Database Design</b>	<b>543</b>
12.11.1	Data Fragmentation	543
12.11.2	Data Replication	547
12.11.3	Data Allocation	548
<b>12.12</b>	<b>The CAP Theorem</b>	<b>549</b>
<b>12.13</b>	<b>Hadoop Distributed File System</b>	<b>550</b>
<b>12.14</b>	<b>C. J. Date's 12 Commandments for Distributed Databases</b>	<b>552</b>
	<b>Summary</b>	<b>554</b>
	<b>Key Terms</b>	<b>555</b>
	<b>Review Questions</b>	<b>555</b>
	<b>Problems</b>	<b>556</b>
<b>CHAPTER 13</b>	<b>BUSINESS INTELLIGENCE AND DATA WAREHOUSES</b>	<b>559</b>
<b>13.1</b>	<b>The Need for Data Analysis</b>	<b>560</b>
<b>13.2</b>	<b>Business Intelligence</b>	<b>560</b>
13.2.1	Business Intelligence Architecture	562
13.2.2	Business Intelligence Benefits	567
13.2.3	Business Intelligence Evolution	567
13.2.4	Business Intelligence Technology Trends	570
<b>13.3</b>	<b>Decision Support Data</b>	<b>571</b>
13.3.1	Operational Data vs. Decision Support Data	571
13.3.2	Decision Support Database Requirements	574
<b>13.4</b>	<b>The Data Warehouse</b>	<b>576</b>
13.4.1	Data Marts	578
13.4.2	Twelve Rules That Define a Data Warehouse	579
<b>13.5</b>	<b>Star Schemas</b>	<b>580</b>
13.5.1	Facts	580
13.5.2	Dimensions	580
13.5.3	Attributes	581
13.5.4	Attribute Hierarchies	582
13.5.5	Star Schema Representation	584
13.5.6	Performance-Improving Techniques for the Star Schema	587
<b>13.6</b>	<b>Data Analytics</b>	<b>590</b>
13.6.1	Data Mining	590
13.6.2	Predictive Analytics	593
<b>13.7</b>	<b>Online Analytical Processing</b>	<b>594</b>
13.7.1	Multidimensional Data Analysis Techniques	594
13.7.2	Advanced Database Support	596
13.7.3	Easy-to-Use End-User Interfaces	596
13.7.4	OLAP Architecture	597
13.7.5	Relational OLAP	599
13.7.6	Multidimensional OLAP	601
13.7.7	Relational vs. Multidimensional OLAP	602
<b>13.8</b>	<b>SQL Extensions for OLAP</b>	<b>603</b>
13.8.1	The ROLLUP Extension	604
13.8.2	The CUBE Extension	605
13.8.3	Materialized Views	606

## TABLE OF CONTENTS

<b>Summary</b>	<b>611</b>
<b>Key Terms</b>	<b>612</b>
<b>Review Questions</b>	<b>612</b>
<b>Problems</b>	<b>613</b>

### PART V DATABASES AND THE INTERNET

<b>Business Vignette</b>	<b>625</b>
<b>CHAPTER 14 DATABASE CONNECTIVITY AND WEB TECHNOLOGIES</b>	<b>626</b>
<b>14.1 Database Connectivity</b>	<b>627</b>
14.1.1 Native SQL Connectivity	627
14.1.2 ODBC, DAO, and RDO	628
14.1.3 OLE-DB	631
14.1.4 ADO.NET	634
14.1.5 Java Database Connectivity (JDBC)	636
<b>14.2 Database Internet Connectivity</b>	<b>638</b>
14.2.1 Web-to-Database Middleware: Server-Side Extensions	639
14.2.2 Web Server Interfaces	641
14.2.3 The Web Browser	642
14.2.4 Client-Side Extensions	643
14.2.5 Web Application Servers	644
14.2.6 Web Database Development	644
<b>14.3 Extensible Markup Language (XML)</b>	<b>648</b>
14.3.1 Document Type Definitions (DTD) and XML Schemas	650
14.3.2 XML Presentation	653
14.3.3 XML Applications	655
<b>14.4 Cloud Computing Services</b>	<b>656</b>
14.4.1 Cloud Implementation Types	658
14.4.2 Characteristics of Cloud Services	659
14.4.3 Types of Cloud Services	660
14.4.4 Cloud Services: Advantages and Disadvantages	661
14.4.5 SQL Data Services	663
<b>Summary</b>	<b>664</b>
<b>Key Terms</b>	<b>665</b>
<b>Review Questions</b>	<b>665</b>
<b>Problems</b>	<b>667</b>

### PART VI DATABASE ADMINISTRATION

<b>Business Vignette</b>	<b>669</b>
<b>CHAPTER 15 DATABASE ADMINISTRATION AND SECURITY</b>	<b>670</b>
<b>15.1 Data as a Corporate Asset</b>	<b>671</b>
<b>15.2 The Need for a Database and Its Role in an Organization</b>	<b>672</b>
<b>15.3 Introduction of a Database: Special Considerations</b>	<b>674</b>

## TABLE OF CONTENTS

<b>15.4 The Evolution of Database Administration</b>	<b>675</b>
<b>15.5 The Database Environment's Human Component</b>	<b>678</b>
15.5.1 The DBA's Managerial Role	680
15.5.2 The DBA's Technical Role	684
15.5.3 The DBA's Role in the Cloud	690
<b>15.6 Security</b>	<b>691</b>
15.6.1 Security Policies	692
15.6.2 Security Vulnerabilities	692
15.6.3 Database Security	694
<b>15.7 Database Administration Tools</b>	<b>695</b>
15.7.1 The Data Dictionary	695
15.7.2 CASE Tools	698
<b>15.8 Developing a Data Administration Strategy</b>	<b>700</b>
<b>15.9 The DBA at Work: Using Oracle for Database Administration</b>	<b>702</b>
15.9.1 Oracle Database Administration Tools	702
15.9.2 Ensuring that the RDBMS Starts Automatically	703
15.9.3 Creating Tablespaces and Datafiles	704
15.9.4 Managing Users and Establishing Security	706
15.9.5 Customizing the Database Initialization Parameters	709
<b>Summary</b>	<b>711</b>
<b>Key Terms</b>	<b>712</b>
<b>Review Questions</b>	<b>712</b>

<b>GLOSSARY</b>	<b>715</b>
<b>INDEX</b>	<b>735</b>

The following appendixes are included on the Instructor and Student Companion Sites or at the CourseMate site for this text, [www.cengagebrain.com](http://www.cengagebrain.com).

<b>APPENDIX A</b>	<b>DESIGNING DATABASES WITH VISIO PROFESSIONAL: A TUTORIAL</b>
<b>APPENDIX B</b>	<b>THE UNIVERSITY LAB: CONCEPTUAL DESIGN</b>
<b>APPENDIX C</b>	<b>THE UNIVERSITY LAB: CONCEPTUAL DESIGN VERIFICATION, LOGICAL DESIGN, AND IMPLEMENTATION</b>
<b>APPENDIX D</b>	<b>CONVERTING AN ER MODEL INTO A DATABASE STRUCTURE</b>
<b>APPENDIX E</b>	<b>COMPARISON OF ER MODEL NOTATIONS</b>
<b>APPENDIX F</b>	<b>CLIENT/SERVER SYSTEMS</b>
<b>APPENDIX G</b>	<b>OBJECT-ORIENTED DATABASES</b>
<b>APPENDIX H</b>	<b>UNIFIED MODELING LANGUAGE (UML)</b>
<b>APPENDIX I</b>	<b>DATABASES IN ELECTRONIC COMMERCE</b>
<b>APPENDIX J</b>	<b>WEB DATABASE DEVELOPMENT WITH COLD FUSION</b>
<b>APPENDIX K</b>	<b>THE HIERARCHICAL DATABASE MODEL</b>
<b>APPENDIX L</b>	<b>THE NETWORK DATABASE MODEL</b>
<b>APPENDIX M</b>	<b>MICROSOFT® ACCESS® TUTORIAL</b>
<b>APPENDIX N</b>	<b>CREATING A NEW DATABASE USING ORACLE 12C</b>
<b>APPENDIX O</b>	<b>DATA WAREHOUSE IMPLEMENTATION FACTORS</b>

The CourseMate Site for this text is available at [www.cengagebrain.com](http://www.cengagebrain.com).

## PREFACE

An eleventh edition is a milestone that few textbooks achieve. We wrote the first edition of this book because we wanted to explain the complexity of database systems in a language that was easy for students to understand. Over the years, we have maintained this emphasis on reaching out to students to explain complex concepts in a practical, approachable manner. This book has been successful through ten editions because the authors, editors, and the publisher paid attention to the impact of technology and to adopter questions and suggestions. We believe that this eleventh edition successfully reflects the same attention to such factors.

In many respects, rewriting a book is more difficult than writing it the first time. If the book is successful, as this one is, a major concern is that the updates, inserts, and deletions will adversely affect writing style and continuity of coverage. The combination of superb reviewers and editors, plus a wealth of feedback from adopters and students of the previous editions, helped make this new edition the best yet.



### CHANGES TO THE ELEVENTH EDITION

In this eleventh edition, we have added some new features and reorganized some coverage to provide a better flow of material. Aside from enhancing the already strong coverage of database design, we have made other improvements in the topical coverage. Here are a few of the highlights:

- Expanded relational algebra coverage with formal definitions and notations
- Updated Business Vignettes showing the impact of database technologies in the real world
- Updated coverage of cloud data services
- Expanded coverage of Big Data and related Hadoop technologies
- Added coverage of data visualization
- SQL coverage expanded to include MySQL databases
- Improved readability and overall visual appeal

This eleventh edition continues to provide a solid and practical foundation for the design, implementation, and management of database systems. This foundation is built on the notion that, while databases are very practical, their successful creation depends on understanding the important concepts that define them. It's not easy to come up with the proper mix of theory and practice, but the previously mentioned feedback suggests that we largely succeeded in our quest to maintain the proper balance.

### THE APPROACH: A CONTINUED EMPHASIS ON DESIGN

As the title suggests, *Database Systems: Design, Implementation, and Management* covers three broad aspects of database systems. However, for several important reasons, special attention is given to database design.

- The availability of excellent database software enables people with little experience to create databases and database applications. Unfortunately, the “create without design” approach usually paves the road to any number of database disasters. In our experience, many database system failures are traceable to poor design and cannot be solved with the help of even the best programmers and managers. Nor is better DBMS software likely to overcome problems created or magnified by poor design. Even the best bricklayers and carpenters can't create a good building from a bad blueprint.
- Most vexing problems of database system management seem to be triggered by poorly designed databases. It hardly seems worthwhile to use scarce resources to develop excellent database management skills merely to use them on crises induced by poorly designed databases.
- Design provides an excellent means of communication. Clients are more likely to get what they need when database system design is approached carefully and thoughtfully. In fact, clients may discover how their organizations really function once a good database design is completed.

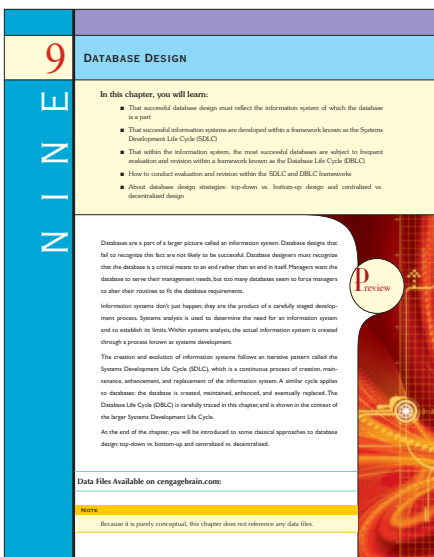
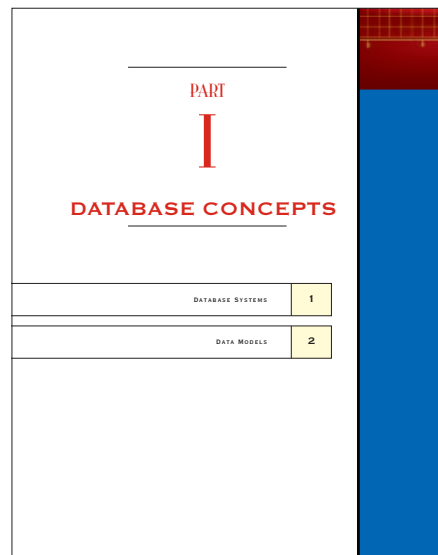
- Familiarity with database design techniques promotes understanding of current database technologies. For example, because data warehouses derive much of their data from operational databases, data warehouse concepts, structures, and procedures make more sense when the operational database's structure and implementation are understood.

Because the practical aspects of database design are stressed, we have covered design concepts and procedures in detail, making sure that the numerous end-of-chapter problems and cases are sufficiently challenging so students can develop real and useful design skills. We also make sure that students understand the potential and actual conflicts between database design elegance, information requirements, and transaction processing speed. For example, it makes little sense to design databases that meet design elegance standards while they fail to meet end-user information requirements. Therefore, we explore the use of carefully defined trade-offs to ensure that the databases meet end-user requirements while conforming to high design standards.

## TOPICAL COVERAGE

### The Systems View

The book's title begins with *Database Systems*. Therefore, we examine the database and design concepts covered in Chapters 1–6 as part of a larger whole by placing them within the systems analysis framework of Chapter 9. Database designers who fail to understand that the database is part of a larger system are likely to overlook important design requirements. In fact, Chapter 9, Database Design, provides the map for the advanced database design developed in Appendixes B and C. Within the larger systems framework, we can also explore issues such as transaction management and concurrency control (Chapter 10), distributed database management systems (Chapter 12), business intelligence and data warehouses (Chapter 13), database connectivity and web technologies (Chapter 14), and database administration and security (Chapter 15).



### Database Design

The first item in the book's subtitle is *Design*, and our examination of database design is comprehensive. For example, Chapters 1 and 2 examine the development and future of databases and data models, and illustrate the need for design. Chapter 3 examines the details of the relational database model; Chapter 4 provides extensive, in-depth, and practical database design coverage; and Chapter 5 explores advanced database design topics. Chapter 6 is devoted to critical normalization issues that affect database efficiency and effectiveness. Chapter 9 examines database design within the systems framework and maps the activities required to successfully design and implement the complex, real-world database developed in Appendixes B and C. Appendix A, *Designing Databases with Visio Professional: A Tutorial*, provides a good introductory tutorial for the use of a database design tool.

Because database design is affected by real-world transactions, the way data are distributed, and ever-increasing information requirements, we examine major database features that must be supported in current-generation databases and models. For example, Chapter 10, Transaction Management and Concurrency Control, focuses on the characteristics of database transactions and how they affect database integrity and consistency. Chapter 11, Database Performance Tuning and Query Optimization, illustrates the need for query efficiency in a world that routinely generates and uses terabyte-sized databases and tables with millions of records. Chapter 12, Distributed Database Management Systems, focuses on data distribution, replication, and allocation. In Chapter 13, Business Intelligence and Data Warehouses, we explore the characteristics of databases that are used in decision support and online analytical processing. Chapter 14, Database Connectivity and Web Technologies, covers the basic database connectivity issues in a web-based data world, development of web-based database front ends, and emerging cloud-based services.

### Implementation

The second portion of the subtitle is **Implementation**. We use Structured Query Language (SQL) in Chapters 7 and 8 to show how databases are implemented and managed. Appendix M, Microsoft Access Tutorial, provides a quick but comprehensive guide to implementing an MS Access database. Appendixes B and C demonstrate the design of a database that was fully implemented; these appendixes illustrate a wide range of implementation issues. We had to deal with conflicting design goals: design elegance, information requirements, and operational speed. Therefore, we carefully audited the initial design in Appendix B to check its ability to meet end-user needs and establish appropriate implementation protocols. The result of this audit yielded the final design developed in Appendix C. The special issues encountered in an Internet database environment are addressed in Chapter 14, Database Connectivity and Web Technologies, and in Appendix J, Web Database Development with ColdFusion.

PART <b>III</b> ADVANCED DESIGN AND IMPLEMENTATION	
INTRODUCTION TO STRUCTURED QUERY LANGUAGE (SQL)	7
ADVANCED SQL	8
DATABASE DESIGN	9

### Management

The final portion of the subtitle is **Management**.

We deal with database management issues in Chapter 10, Transaction Management and Concurrency Control; Chapter 12, Distributed Database Management Systems; and Chapter 15, Database Administration and Security. Chapter 11, Database Performance Tuning and Query Optimization, is a valuable resource that illustrates how a DBMS manages data retrieval. In addition, Appendix N, Creating a New Database Using Oracle 12c, walks you through the process of setting up a new database.

PART <b>VI</b> DATABASE ADMINISTRATION	
DATABASE ADMINISTRATION AND SECURITY	15


TEACHING DATABASE: A MATTER OF FOCUS

Given the wealth of detailed coverage, instructors can “mix and match” chapters to produce the desired coverage. Depending on where database courses fit into the curriculum, instructors may choose to emphasize database design or database management. (See Figure 1.)

The hands-on nature of database design lends itself particularly well to class projects for which students use instructor-selected software to prototype a system they design for the end user. Several end-of-chapter problems are sufficiently complex to serve as projects, or an instructor may work with local businesses to give students hands-on experience. Note that some elements of the database design track are also found in the database management track, because it is difficult to manage database technologies that are not well understood.

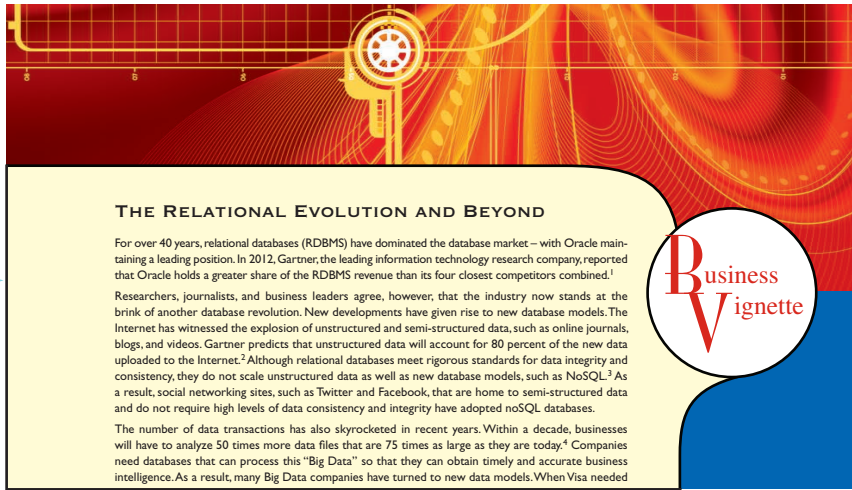
The options shown in Figure 1 serve only as a starting point. Naturally, instructors will tailor their coverage based on their specific course requirements. For example, an instructor may decide to make Appendix I an outside reading assignment and Appendix A a self-taught tutorial, and then use that time to cover client/server systems or object-oriented databases. The latter choice would serve as a gateway to UML coverage.

FIGURE 1



**Business Vignettes**

highlight topics in a real-life setting.



**THE RELATIONAL EVOLUTION AND BEYOND**

For over 40 years, relational databases (RDBMS) have dominated the database market – with Oracle maintaining a leading position. In 2012, Gartner, the leading information technology research company, reported that Oracle holds a greater share of the RDBMS revenue than its four closest competitors combined.<sup>1</sup>

Researchers, journalists, and business leaders agree, however, that the industry now stands at the brink of another database revolution. New developments have given rise to new database models. The Internet has witnessed the explosion of unstructured and semi-structured data, such as online journals, blogs, and videos. Gartner predicts that unstructured data will account for 80 percent of the new data uploaded to the Internet.<sup>2</sup> Although relational databases meet rigorous standards for data integrity and consistency, they do not scale unstructured data as well as new database models, such as NoSQL.<sup>3</sup> As a result, social networking sites, such as Twitter and Facebook, that are home to semi-structured data and do not require high levels of data consistency and integrity have adopted noSQL databases.

The number of data transactions has also skyrocketed in recent years. Within a decade, businesses will have to analyze 50 times more data files that are 75 times as large as they are today.<sup>4</sup> Companies need databases that can process this “Big Data” so that they can obtain timely and accurate business intelligence. As a result, many Big Data companies have turned to new data models. When Visa needed

**Business Vignette**

**Online Content boxes**

draw attention to material at [www.cengagebrain.com](http://www.cengagebrain.com) for this text and provide ideas for incorporating this content into the course.



**ONLINE CONTENT**

The databases used in each chapter are available at [www.cengagebrain.com](http://www.cengagebrain.com). Throughout the book, Online Content boxes highlight material related to content at [www.cengagebrain.com](http://www.cengagebrain.com).

**Notes** highlight important facts about the concepts introduced in the chapter.


**NOTE**

Data that display data inconsistency are also referred to as data that lack data integrity. **Data integrity** is defined as the condition in which all of the data in the database are consistent with the real-world events and conditions. In other words, data integrity means that:

- Data are *accurate*—there are no data inconsistencies.
- Data are *verifiable*—the data will always yield consistent results.

A variety of **four-color figures**, including ER models and implementations, tables, and illustrations, clearly illustrate difficult concepts.

**FIGURE 1.10** Illustrating metadata with Microsoft SQL Server Express





## TEXT FEATURES

### S U M M A R Y

- The ERM uses ERDs to represent the conceptual database as viewed by the end user. The ERM's main components are entities, relationships, and attributes. The ERD also includes connectivity and cardinality notations. An ERD can also show relationship strength, relationship participation (optional or mandatory), and degree of relationship (unary, binary, ternary, etc.).
- Connectivity describes the relationship classification (1:1, 1:M, or M:N). Cardinality expresses the specific number of entity occurrences associated with an occurrence of a related entity. Connectivities and cardinalities are usually based on business rules.
- In the ERM, an M:N relationship is valid at the conceptual level. However, when implementing the ERM in a relational database, the M:N relationship must be mapped to a set of 1:M relationships through a composite entity.

A robust **Summary** at the end of each chapter ties together the major concepts and serves as a quick review for students.

### K E Y T E R M S

algorithms	data mart	explanatory analytics
attribute hierarchy	data mining	extraction, transformation, and load-
business intelligence (BI)	data warehouse	ing (ETL)
cube cache	decision support system (DSS)	fact table
dashboard	dimension tables	facts
data analytics	dimensions	governance
data cube	drill down	key performance indicators (KPIs)

An alphabetic list of **Key Terms** summarizes important terms.

### R E V I E W Q U E S T I O N S

1. What two conditions must be met before an entity can be classified as a weak entity? Give an example of a weak entity.
2. What is a strong (or identifying) relationship, and how is it depicted in a Crow's Foot ERD?
3. Given the business rule "an employee may have many degrees," discuss its effect on attributes, entities, and relationships. (*Hint*: Remember what a multivalued attribute is and how it might be implemented.)
4. What is a composite entity, and when is it used?
5. Suppose you are working within the framework of the conceptual model in Figure Q4.5.

**Review Questions** challenge students to apply the skills learned in each chapter.

### P R O B L E M S

1. Given the following business rules, create the appropriate Crow's Foot ERD.
  - a. A company operates many departments.
  - b. Each department employs one or more employees.
  - c. Each of the employees may or may not have one or more dependents.
  - d. Each employee may or may not have an employment history.
2. The Hudson Engineering Group (HEG) has contacted you to create a conceptual model whose application will meet the expected database requirements for the company's training program. The HEG administrator gives you

**Problems** become progressively more complex as students draw on the lessons learned from the completion of preceding problems.

### COURSEMATE

Cengage Learning's *Database Systems CourseMate* brings course concepts to life with interactive learning, study, and exam preparation tools that support the printed textbook. Watch student comprehension soar as your class works with the printed textbook and the textbook-specific website. CourseMate goes beyond the book to deliver what you need! Learn more at [www.cengage.com/coursemate](http://www.cengage.com/coursemate).

#### Engagement Tracker

How do you assess your students' engagement in your course? How do you know students have read the material or viewed the resources you've assigned? How can you tell if your students are struggling with a concept? With CourseMate, you can use the included Engagement Tracker to assess student preparation and engagement. Use the tracking tools to see progress for the class as a whole or for individual students. Identify students at risk early in the course. Uncover which concepts are most difficult for your class. Monitor time on task. Keep your students engaged.

#### Interactive Teaching and Learning Tools

CourseMate includes interactive teaching and learning tools:

- Quizzes
- Crossword puzzles
- Flashcards
- Videos
- PowerPoint presentations
- and more

These assets enable students to review for tests, prepare for class, and address the needs of their varied learning styles.

#### Interactive eBook

MindTap Reader is Cengage Learning's re-imagining of the traditional eBook, specifically designed for assimilating content and media assets in a fully online – and often mobile – reading environment. MindTap Reader combines thoughtful navigation ergonomics, advanced annotation support and interactivity through the placement of inline documents and media assets. These capabilities create an engaging reading experience further enhanced through tightly integrated web-apps (e.g. text-to-speech, note-taking, utilities) delivering a learning tool that drives immediacy, relevancy and engagement for today's learners.

#### Appendixes

Fifteen online appendixes provide additional material on a variety of important areas, such as using Microsoft® Visio® and Microsoft® Access®, ER model notations, UML, object-oriented databases, databases and electronic commerce, and Adobe® ColdFusion®.

#### Database, SQL Script, and ColdFusion Files

The online materials for this book include all of the database structures and table contents used in the text. For students using Oracle®, MySQL and Microsoft SQL Server™, SQL scripts are included to help students create and load all tables used in the SQL chapters (7 and 8). In addition, all ColdFusion scripts used to develop the web interfaces in Appendix J are included.

## ADDITIONAL FEATURES

### INSTRUCTOR RESOURCES

*Database Systems: Design, Implementation, and Management, Eleventh Edition*, includes teaching tools to support instructors in the classroom. The ancillary material that accompanies the textbook is listed below. They are available on the web at [www.cengage.com](http://www.cengage.com).

#### **Instructor's Manual**

The authors have created this manual to help instructors make their classes informative and interesting. Because the authors tackle so many problems in depth, instructors will find the *Instructor's Manual* especially useful. The details of the design solution process are shown in the *Instructor's Manual*, as well as notes about alternative approaches that may be used to solve a particular problem.

#### **SQL Script Files for Instructors**

The authors have provided teacher's SQL script files to let instructors cut and paste the SQL code into the SQL windows. (Scripts are provided for Oracle, MySQL, and MS SQL Server.) The SQL scripts, which have all been tested by Cengage Learning, are a major convenience for instructors. You won't have to type in the SQL commands, and the use of the scripts eliminates typographical errors that are sometimes difficult to trace.

#### **ColdFusion Files for Instructors**

The ColdFusion web development solutions are provided. Instructors have access to a menu-driven system that lets teachers show the code as well as its execution.

#### **Databases**

For many chapters, Microsoft® Access® instructor databases are available that include features not found in the student databases. For example, the databases that accompany Chapters 7 and 8 include many of the queries that produce the problem solutions. Other Access databases, such as the ones that accompany Chapters 3, 4, 5, and 6, include implementations of the design problem solutions to let instructors illustrate the effect of design decisions. In addition, instructors have access to all the script files for Oracle, MySQL, and MS SQL Server so that all the databases and their tables can be converted easily and precisely.

#### **Cengage Learning Testing Powered by Cognero**

##### **A flexible, online system that allows you to:**

- Author, edit, and manage test bank content from multiple Cengage Learning solutions
- Create multiple test versions in an instant
- Deliver tests from your LMS, your classroom or wherever you want

##### **Start right away!**

Cengage Learning Testing Powered by Cognero works on any operating system or browser.

- No special installs or downloads needed
- Create tests from school, home, the coffee shop – anywhere with Internet access

### What will you find?

- Simplicity at every step. A desktop-inspired interface features drop-down menus and familiar, intuitive tools that take you through content creation and management with ease.
- Full-featured test generator. Create ideal assessments with your choice of 15 question types (including true/false, multiple choice, opinion scale/likert, and essay). Multi-language support, an equation editor and unlimited metadata help ensure your tests are complete and compliant.
- Cross-compatible capability. Import and export content into other systems.

### PowerPoint® Presentations

Microsoft PowerPoint slides are included for each chapter. Instructors can use the slides in a variety of ways—for example, as teaching aids during classroom presentations or as printed handouts for classroom distribution. Instructors can modify these slides or include slides of their own for additional topics introduced to the class.

### Figure Files

Figure files for solutions are presented in the Instructor's Manual to allow instructors to create their own presentations. Instructors can also manipulate these files to meet their particular needs.

## ACKNOWLEDGMENTS

Regardless of how many editions of this book are published, they will always rest on the solid foundation created by the first edition. We remain convinced that our work has become successful because that first edition was guided by Frank Ruggirello, a former Wadsworth senior editor and publisher. Aside from guiding the book's development, Frank also managed to solicit the great Peter Keen's evaluation (thankfully favorable) and subsequently convinced PK to write the foreword for the first edition. Although we sometimes found Frank to be an especially demanding taskmaster, we also found him to be a superb professional and a fine friend. We suspect Frank will still see his fingerprints all over our current work. Many thanks.

A difficult task in rewriting a book is deciding what new approaches, topical coverage, and changes to depth of coverage are appropriate for a product that has successfully weathered the test of the marketplace. The comments and suggestions made by the book's adopters, students, and reviewers play a major role in deciding what coverage is desirable and how that coverage is to be treated.

Some adopters became extraordinary reviewers, providing incredibly detailed and well-reasoned critiques even as they praised the book's coverage and style. Dr. David Hatherly, a superb database professional who is a senior lecturer in the School of Information Technology, Charles Sturt University–Mitchell, Bathurst, Australia, made sure that we knew precisely what issues led to his critiques. Even better for us, he provided the suggestions that made it much easier for us to improve the topical coverage in earlier editions. Dr. Hatherly's recommendations continue to be reflected in this tenth edition. All of his help was given freely and without prompting on our part. His efforts are much appreciated, and our thanks are heartfelt.

We also owe a debt of gratitude to Professor Emil T. Cipolla, who teaches at St. Mary College. Professor Cipolla's wealth of IBM experience turned out to be a valuable resource when we tackled the embedded SQL coverage in Chapter 8.

Every technical book receives careful scrutiny by several groups of reviewers selected by the publisher. We were fortunate to face the scrutiny of reviewers who were superbly qualified to offer their critiques, comments, and suggestions—many of which strengthened this edition. While holding them blameless for any remaining shortcomings, we owe these reviewers many thanks for their contributions:

Mubarak Banisaklher, Bethune Cookman University

David Bell, Pacific Union College

Yurii Boreisha, Minnesota State University Moorhead

Laurie Crawford, Franklin University

Mel Goetting, Shawnee State University

Jeff Guan, University of Louisville

William Hochstettler, Franklin University

Laurene Hutchinson, Louisiana State University,  
Baton Rouge

Nitin Kale, University of Southern California, Los Angeles

Gerald Karush, Southern New Hampshire University

Michael Kelly, Community College of Rhode Island

Timothy Koets, Grand Rapids Community College

Klara Nelson, The University of Tampa

Chiso Okafor, Roxbury Community College

Brandon Olson, The College of St. Scholastica

James Reneau, Shawnee State University

Julio Rivera, University of Alabama at Birmingham

Ruth Robins, University of Houston Downtown

Samuel Sambasivam, Azusa Pacific University

Paul Seibert, North Greenville University

Ronghua Shan, Dakota State University

Andrew Smith, Marian University

Antonis Stylianou, University of North Carolina Charlotte

Brian West, University of Louisiana at Lafayette

Nathan White, McKendree University



## ACKNOWLEDGMENTS

In some respects, writing books resembles building construction: When 90 percent of the work seems done, 90 percent of the work remains to be done. Fortunately for us, we had a great team on our side.

- We are deeply indebted to Deb Kaufmann for her help and guidance. Deb has been everything we could have hoped for in a development editor and more. She brought fresh eyes and a new perspective that meshed wonderfully with our vision for the new edition.
- After writing so many books and ten editions of *this* book, we know just how difficult it can be to transform the authors' work into an attractive product. The production team, both at Cengage Learning (Matt Hutchinson) and GEX Publishing Services (Marisa Taylor), have done an excellent job.
- We also owe Kate Mason, our Content Developer, special thanks for her ability to guide this book to a successful conclusion. Kate's work touched on all of the publication bases, and her managerial skills protected us from publishing gremlins that might have become a major nuisance. Also, her skills in dealing with occasionally cranky authors far exceed those of any diplomat we can think of. And did we mention that Kate is, quite simply, a delightful person? We would also like to thank Anne Merrill, Content Developer, who competently shepherded the text in its later stages.

We also thank our students for their comments and suggestions. They are the reason for writing this book in the first place. One comment stands out in particular: "I majored in systems for four years, and I finally discovered why when I took your course." And one of our favorite comments by a former student was triggered by a question about the challenges created by a real-world information systems job: "Doc, it's just like class, only easier. You really prepared me well. Thanks!"

Special thanks go to a very unique and charismatic gentleman. For over 20 years Peter Rob has been the driving force behind the creation and evolution of this book. This book originated as a product of his drive and dedication to excellence. For over twenty years, he was the voice of *Database Systems*, and the driving force behind its advancement. We wish him peace in his retirement, time with his loved ones, and luck on his many projects.

Last, and certainly not least, we thank our families for their solid support at home. They graciously accepted the fact that during more than a year's worth of rewriting, there would be no free weekends, rare free nights, and even rarer free days. We owe you much, and the dedications we wrote are but a small reflection of the important space you occupy in our hearts.

Carlos Coronel and Steven Morris

---

PART

I

**DATABASE CONCEPTS**

---

DATABASE SYSTEMS	1
DATA MODELS	2

## THE RELATIONAL EVOLUTION AND BEYOND

For over 40 years, relational databases (RDBMS) have dominated the database market – with Oracle maintaining a leading position. In 2012, Gartner, the leading information technology research company, reported that Oracle holds a greater share of the RDBMS revenue than its four closest competitors combined.<sup>1</sup>

Researchers, journalists, and business leaders agree, however, that the industry now stands at the brink of another database revolution. New developments have given rise to new database models. The Internet has witnessed the explosion of unstructured and semi-structured data, such as online journals, blogs, and videos. Gartner predicts that unstructured data will account for 80 percent of the new data uploaded to the Internet.<sup>2</sup> Although relational databases meet rigorous standards for data integrity and consistency, they do not scale unstructured data as well as new database models, such as NoSQL.<sup>3</sup> As a result, social networking sites, such as Twitter and Facebook, that are home to semi-structured data and do not require high levels of data consistency and integrity have adopted noSQL databases.

The number of data transactions has also skyrocketed in recent years. Within a decade, businesses will have to analyze 50 times more data files that are 75 times as large as they are today.<sup>4</sup> Companies need databases that can process this “Big Data” so that they can obtain timely and accurate business intelligence. As a result, many Big Data companies have turned to new data models. When Visa needed to process two years of credit transactions, the company turned to Hadoop to process 70 billion transactions, thereby reducing processing time from one month to 13 minutes.<sup>5</sup> Another technology that is fast gaining ground is in-memory databases, which store data in main memory rather than secondary storage. In-memory databases take advantage of parallel computing and multicore processors to deliver faster business intelligence.<sup>6</sup>

SAP’s market share grew by an impressive 48 percent in 2012 (compared to Oracle’s 18 percent) and the company attributes this growth to the SAP HANNA In-Memory Appliance.<sup>7</sup>

Additionally, cloud computing provides on-demand access to clusters of computers over the Internet at a reasonable cost. Cloud computing enables the widespread adoption of powerful databases in organizations that could not previously afford the high cost of massive server infrastructure.<sup>8</sup>

All these technology advances are revolutionizing the database market. However, the relational model is not standing idly by. Today, most large relational database vendors provide in-memory database and cloud capabilities. Oracle has invested heavily in research and has acquired many smaller companies in order to expand its capabilities. At the end of 2011, it acquired RightNow Technologies for \$1.5 billion. The deal gave Oracle a leading position in cloud customer service.<sup>9</sup>

The relational model has a history of overcoming challenges from other emerging database technologies, which is why the overwhelming majority of existing systems still rely on relational databases. It remains to be seen if the relational model will evolve once again and adapt to these new challenges. The race is on.

<sup>1</sup> Mark Fontecchio, “Oracle the clear leader in \$24 billion RDBMS market,” *IT Knowledge Exchange*, April 12, 2012, <http://itknowledgeexchange.techtarget.com/eye-on-oracle/oracle-the-clear-leader-in-24-billion-rdbms-market/>.

<sup>2</sup> Sameer Nori, “Are You Ready for the Growth in Unstructured Information?” *ATTIVO Active Intelligence*, March 7, 2013, <http://www.attivio.com/blog/55-industry-insights/1217-are-you-ready-for-the-growth-in-unstructured-information.html>.

<sup>3</sup> Guy Harrison, “10 things you should know about NoSQL databases,” *TechRepublic*, August 26, 2010, <http://www.techrepublic.com/blog/10things/10-things-you-should-know-about-nosql-databases/1772>.

<sup>4</sup> David Rosenbaum, “That New Big Data Magic,” *CFO*, August 29, 2011, [http://www3.cfo.com/article/2011/8/analytics\\_that-new-big-data-magic](http://www3.cfo.com/article/2011/8/analytics_that-new-big-data-magic).

<sup>5</sup> David Strom, “NoSQL: Breaking free of structured data,” *IT World*, June 9, 2011, <http://www.itworld.com/data-centerservers/172477/nosql-breaking-free-structured-data>.

<sup>6</sup> Mike Simons, “SAP Cofounder Hasso Plattner Promises in-Memory Database Revolution,” *CIO*, May 14, 2009, [http://www.cio.com/article/492832/SAP\\_Cofounder\\_Hasso\\_Plattner\\_Promises\\_in\\_Memory\\_Database\\_Revolution](http://www.cio.com/article/492832/SAP_Cofounder_Hasso_Plattner_Promises_in_Memory_Database_Revolution).

<sup>7</sup> Mark Fontecchio, “Oracle the clear leader in \$24 billion RDBMS market,” *IT Knowledge Exchange*, April 12, 2012, <http://itknowledgeexchange.techtarget.com/eye-on-oracle/oracle-the-clear-leader-in-24-billion-rdbms-market/>.

<sup>8</sup> Joshua Greenbaum, “A Revolution Threatens the Relational Database,” *IT Management*, May 13, 2008, <http://itmanagement.earthweb.com/columns/entad/article.php/3746191>.

<sup>9</sup> Larry Dignam, “Oracle acquires RightNow for \$1.5 billion, aims turrets at Salesforce.com,” *ZDNet*, October 24, 2011, <http://www.zdnet.com/blog/bt/oracle-acquires-rightnow-for-1-5-billion-aims-turrets-at-salesforce-com/61681>.

The logo for Business Vignette is located in a white circular bubble on the right side of the page. The word "Business" is written in a red serif font, and "Vignette" is written in a larger, red serif font below it.

**In this chapter, you will learn:**

- The difference between data and information.
- What a database is, the various types of databases, and why they are valuable assets for decision making.
- The importance of database design.
- How modern databases evolved from file systems.
- About flaws in file system data management.
- The main components of the database system.
- The main functions of a database management system (DBMS).

Good decisions require good information that is derived from raw facts. These raw facts are known as data. Data are likely to be managed most efficiently when they are stored in a database. In this chapter, you will learn what a database is, what it does, and why it yields better results than other data management methods. You will also learn about various types of databases and why database design is so important.

Databases evolved from computer file systems. Although file system data management is now largely outmoded, understanding the characteristics of file systems is important because file systems are the source of serious data management limitations. In this chapter, you will also learn how the database system approach helps eliminate most of the shortcomings of file system data management.



**P**review

**Data Files Available on cengagebrain.com:**

Chapter	Available formats				Questions and Problems	Available formats			
	MS Access	Oracle	MS SQL	MySQL		MS Access	Oracle	MS SQL	MySQL
Ch01_Text	✓	✓	✓	✓	Ch01_Problems	✓	✓	✓	✓
Ch01_Design_Example	✓	✓	✓	✓					

## 1.1 WHY DATABASES?

Imagine trying to operate a business without knowing who your customers are, what products you are selling, who is working for you, who owes you money, and to whom you owe money. All businesses have to keep this type of data and much more; just as importantly, they must have those data available to decision makers when necessary. It can be argued that the ultimate purpose of all business information systems is to help businesses use information as an organizational resource. At the heart of all of these systems are the collection, storage, aggregation, manipulation, dissemination, and management of data.

Depending on the type of information system and the characteristics of the business, these data could vary from a few megabytes on just one or two topics to terabytes covering hundreds of topics within the business's internal and external environment. Telecommunications companies, such as Sprint and AT&T, are known to have systems that keep data on trillions of phone calls, with new data being added to the system at speeds up to 70,000 calls per second!<sup>1</sup> Not only do these companies have to store and manage immense collections of data, they have to be able to find any given fact in that data quickly. Consider the case of Internet search staple Google. While Google is reluctant to disclose many details about its data storage specifications, it is estimated that the company responds to over 91 million searches per day across a collection of data that is several terabytes in size. Impressively, the results of these searches are available almost instantly.

How can these businesses process this much data? How can they store it all, and then quickly retrieve just the facts that decision makers want to know, just when they want to know it? The answer is that they use databases. Databases, as explained in detail throughout this book, are specialized structures that allow computer-based systems to store, manage, and retrieve data very quickly. Virtually all modern business systems rely on databases; therefore, a good understanding of how these structures are created and their proper use is vital for any information systems professional. Even if your career does not take you down the amazing path of database design and development, databases will be a key component underpinning the systems that you use. In any case, you will probably make decisions in your career based on information generated from data. Thus, it is important that you know the difference between data and information.

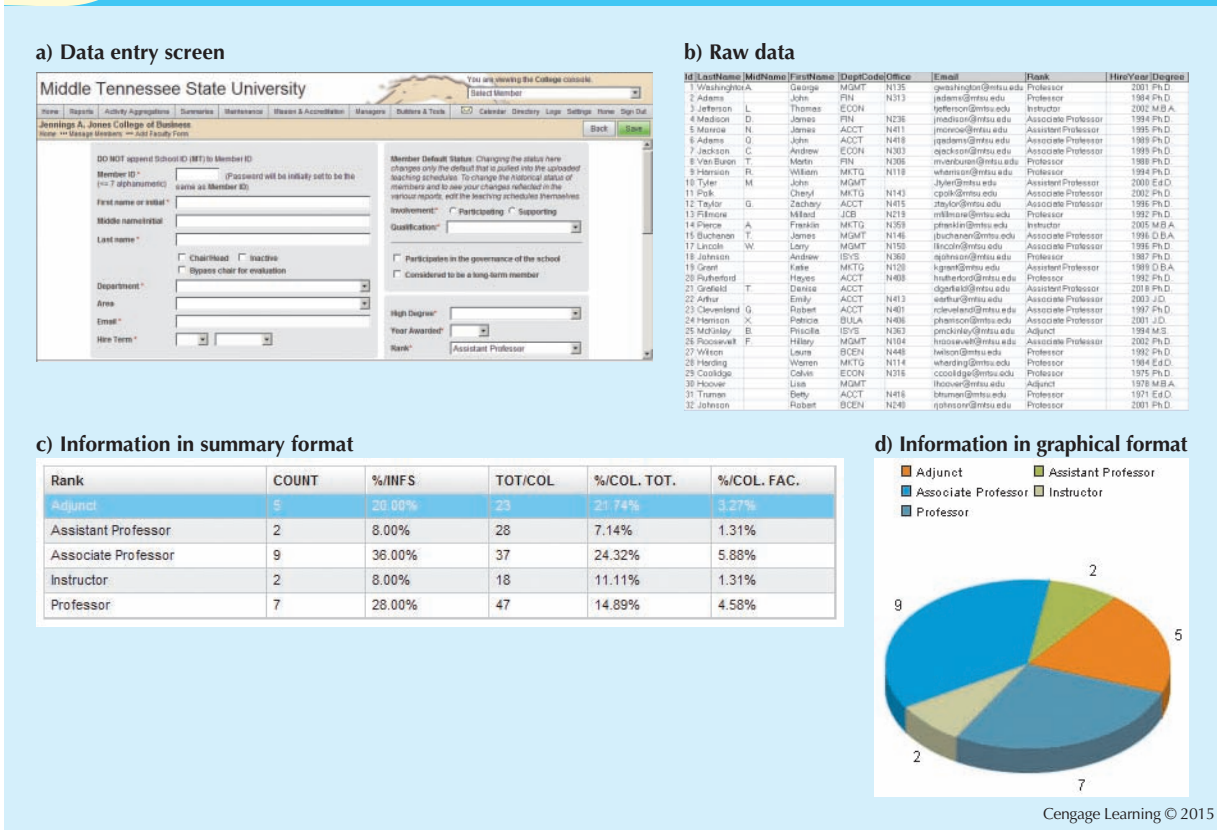
## 1.2 DATA VS. INFORMATION

To understand what drives database design, you must understand the difference between data and information. **Data** are raw facts. The word raw indicates that the facts have not yet been processed to reveal their meaning. For example, suppose that a university tracks data on faculty members for reporting to accrediting bodies. To get the data for each faculty member into the database, you would provide a screen to allow for convenient data entry, complete with drop-down lists, combo boxes, option buttons, and other data-entry validation controls. Figure 1.1(a) shows a simple data-entry form from a software package named Sedona. When the data are entered into the form and saved, they are placed in the underlying database as raw data, as shown in Figure 1.1(b). Although you now have the facts in hand, they are not particularly useful in this format. Reading through hundreds of rows of data for faculty members does not provide much insight into the overall makeup of the faculty. Therefore, you transform the raw data into a data summary like the one shown in Figure 1.1(c). Now you can get quick answers to questions such as “What percentage of the faculty in the Information Systems (INFS) department are adjuncts?” In this case, you can quickly determine that 20 percent of the INFS faculty members are adjunct faculty. Because graphics can enhance your ability to quickly extract meaning from data, you show the data summary pie chart in Figure 1.1(d).

<sup>1</sup> “Top Ten Largest Databases in the World,” *Business Intelligence Lowdown*, February 15, 2007, [http://www.businessintelligencelowdown.com/2007/02/top\\_10\\_largest\\_.html](http://www.businessintelligencelowdown.com/2007/02/top_10_largest_.html).



**FIGURE 1.1** Transforming raw data into information



**Information** is the result of processing raw data to reveal its meaning. Data processing can be as simple as organizing data to reveal patterns or as complex as making forecasts or drawing inferences using statistical modeling. To reveal meaning, information requires *context*. For example, an average temperature reading of 105 degrees does not mean much unless you also know its context: Is this reading in degrees Fahrenheit or Celsius? Is this a machine temperature, a body temperature, or an outside air temperature? Information can be used as the foundation for decision making. For example, the data summary for the faculty can provide accrediting bodies with insights that are useful in determining whether to renew accreditation for the university.

Keep in mind that raw data must be properly *formatted* for storage, processing, and presentation. For example, dates might be stored in Julian calendar formats within the database, but displayed in a variety of formats, such as day-month-year or month/day/year, for different purposes. Respondents' yes/no responses might need to be converted to a Y/N or 0/1 format for data storage. More complex formatting is required when working with complex data types, such as sounds, videos, or images.

In this “information age,” production of accurate, relevant, and timely information is the key to good decision making. In turn, good decision making is the key to business survival in a global market. We are now said to be entering the “knowledge age.”<sup>2</sup> Data are the foundation of information, which is the bedrock of **knowledge**—that is, the body of information and facts about a specific subject. Knowledge implies familiarity, awareness, and understanding of information as it applies to an environment. A key characteristic of knowledge is that “new” knowledge can be derived from “old” knowledge.

<sup>2</sup> Peter Drucker coined the phrase “knowledge worker” in 1959 in his book *Landmarks of Tomorrow*. In 1994, Esther Dyson, George Keyworth, and Dr. Alvin Toffler introduced the concept of the “knowledge age.”

Let's summarize some key points:

- Data constitute the building blocks of information.
- Information is produced by processing data.
- Information is used to reveal the meaning of data.
- Accurate, relevant, and timely information is the key to good decision making.
- Good decision making is the key to organizational survival in a global environment.

Timely and useful information requires accurate data. Such data must be properly generated and stored in a format that is easy to access and process. And, like any basic resource, the data environment must be managed carefully.

**Data management** is a discipline that focuses on the proper generation, storage, and retrieval of data. Given the crucial role that data play, it should not surprise you that data management is a core activity for any business, government agency, service organization, or charity.

## 1.3 INTRODUCING THE DATABASE

Efficient data management typically requires the use of a computer database. A **database** is a shared, integrated computer structure that stores a collection of the following:

- End-user data—that is, raw facts of interest to the end user.
- **Metadata**, or data about data, through which the end-user data are integrated and managed.

The metadata describe the data characteristics and the set of relationships that links the data found within the database. For example, the metadata component stores information such as the name of each data element, the type of values (numeric, dates, or text) stored on each data element, and whether the data element can be left empty. The metadata provide information that complements and expands the value and use of the data. In short, metadata present a more complete picture of the data in the database. Given the characteristics of metadata, you might hear a database described as a “collection of *self-describing* data.”

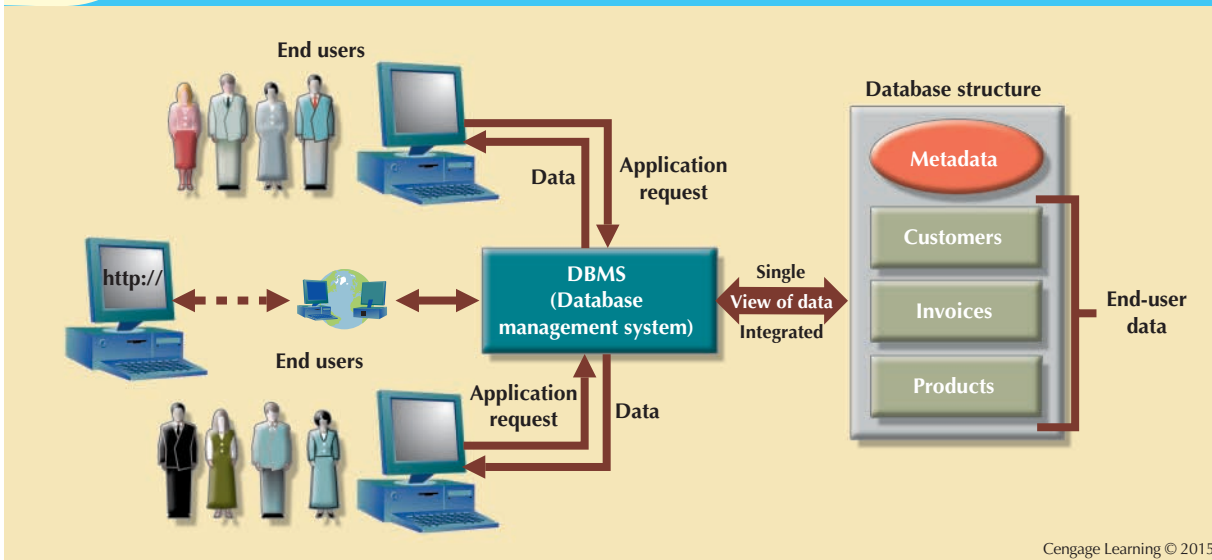
A **database management system (DBMS)** is a collection of programs that manages the database structure and controls access to the data stored in the database. In a sense, a database resembles a very well-organized electronic filing cabinet in which powerful software (the DBMS) helps manage the cabinet's contents.

### 1.3.1 ROLE AND ADVANTAGES OF THE DBMS

The DBMS serves as the intermediary between the user and the database. The database structure itself is stored as a collection of files, and the only way to access the data in those files is through the DBMS. Figure 1.2 emphasizes the point that the DBMS presents the end user (or application program) with a single, integrated view of the data in the database. The DBMS receives all application requests and translates them into the complex operations required to fulfill those requests. The DBMS hides much of the database's internal complexity from the application programs and users. The application program might be written by a programmer using a programming language, such as Visual Basic.NET, Java, or C#, or it might be created through a DBMS utility program.



**FIGURE 1.2** The DBMS manages the interaction between the end user and the database



Having a DBMS between the end user's applications and the database offers some important advantages. First, the DBMS enables the data in the database *to be shared* among multiple applications or users. Second, the DBMS *integrates* the many different users' views of the data into a single all-encompassing data repository.

Because data are the crucial raw material from which information is derived, you must have a good method to manage such data. As you will discover in this book, the DBMS helps make data management more efficient and effective. In particular, a DBMS provides these advantages:

- *Improved data sharing.* The DBMS helps create an environment in which end users have better access to more and better-managed data. Such access makes it possible for end users to respond quickly to changes in their environment.
- *Improved data security.* The more users access the data, the greater the risks of data security breaches. Corporations invest considerable amounts of time, effort, and money to ensure that corporate data are used properly. A DBMS provides a framework for better enforcement of data privacy and security policies.
- *Better data integration.* Wider access to well-managed data promotes an integrated view of the organization's operations and a clearer view of the big picture. It becomes much easier to see how actions in one segment of the company affect other segments.
- *Minimized data inconsistency.* **Data inconsistency** exists when different versions of the same data appear in different places. For example, data inconsistency exists when a company's sales department stores a sales representative's name as Bill Brown and the company's personnel department stores that same person's name as *William G. Brown*, or when the company's regional sales office shows the price of a product as \$45.95 and its national sales office shows the same product's price as \$43.95. The probability of data inconsistency is greatly reduced in a properly designed database.
- *Improved data access.* The DBMS makes it possible to produce quick answers to ad hoc queries. From a database perspective, a **query** is a specific request issued to the DBMS for data manipulation—for example, to read or update the data. Simply put, a query is a question, and an **ad hoc query** is a spur-of-the-moment question. The DBMS sends back an answer (called the **query result set**) to the application. For example, when dealing with large amounts of sales data, end users might want quick answers to questions (ad hoc queries). Some examples include the following:
  - What was the dollar volume of sales by product during the past six months?

- What is the sales bonus figure for each of our salespeople during the past three months?
- How many of our customers have credit balances of \$3,000 or more?
- *Improved decision making.* Better-managed data and improved data access make it possible to generate better-quality information, on which better decisions are based. The quality of the information generated depends on the quality of the underlying data. **Data quality** is a comprehensive approach to promoting the accuracy, validity, and timeliness of the data. While the DBMS does not guarantee data quality, it provides a framework to facilitate data quality initiatives. Data quality concepts will be covered in more detail in Chapter 15, Database Administration and Security.
- *Increased end-user productivity.* The availability of data, combined with the tools that transform data into usable information, empowers end users to make quick, informed decisions that can make the difference between success and failure in the global economy.

The advantages of using a DBMS are not limited to the few just listed. In fact, you will discover many more advantages as you learn more about the technical details of databases and their proper design.

### 1.3.2 TYPES OF DATABASES

A DBMS can be used to build many different types of databases. Each database stores a particular collection of data and is used for a specific purpose. Over the years, as technology and innovative uses of databases have evolved, different methods have been used to classify databases. For example, databases can be classified by the number of users supported, where the data are located, the type of data stored, the intended data usage, and the degree to which the data are structured.

The number of users determines whether the database is classified as single user or multiuser. A **single-user database** supports only one user at a time. In other words, if user A is using the database, users B and C must wait until user A is done. A single-user database that runs on a personal computer is called a **desktop database**. In contrast, a **multiuser database** supports multiple users at the same time. When the multiuser database supports a relatively small number of users (usually fewer than 50) or a specific department within an organization, it is called a **workgroup database**. When the database is used by the entire organization and supports many users (more than 50, usually hundreds) across many departments, the database is known as an **enterprise database**.

Location might also be used to classify the database. For example, a database that supports data located at a single site is called a **centralized database**. A database that supports data distributed across several different sites is called a **distributed database**. The extent to which a database can be distributed and the way in which such distribution is managed are addressed in detail in Chapter 12, Distributed Database Management Systems.

Both centralized and decentralized (distributed) databases require a well-defined infrastructure (hardware, operating systems, network technologies, etc.) to implement and operate the database. Typically, the infrastructure is owned and maintained by the organization that creates and operates the database. But in recent years the use of cloud databases has been growing in popularity. A **cloud database** is a database that is created and maintained using cloud data services, such as Microsoft Azure or Amazon's AWS. These services, provided by third-party vendors, provide defined performance measures (data storage capacity, required throughput, and availability) for the database, but do not necessarily specify the underlying infrastructure to implement it. The data owner does not have to know, or be concerned about, what hardware and software is being used to support their database. The performance capabilities can be renegotiated with the cloud provider as the business demands on the database change. For example, during the 2012 presidential election in the United States, the Obama campaign used a cloud database hosted on infrastructure capabilities purchased from Amazon. The campaign did not have to buy, install, configure, or maintain any hardware, operating systems, or network devices. It simply purchased storage and processing capacity for its data and applications. As the demands on the database increased, additional processing and storage capabilities could be purchased as needed.

In some contexts, such as research environments, a popular way of classifying databases is according to the type of data stored in them. Using this criterion, databases are grouped into two categories: general-purpose and discipline-specific databases. **General-purpose databases** contain a wide variety of data used in multiple disciplines—for example,