



Oracle[®] 12c:
SQL

Joan Casteel

5 REASONS

to buy your textbooks
and course materials at

CENGAGE **brain**.com

- 1 SAVINGS:**
Prices up to 75% off, daily coupons, and free shipping on orders over \$25
- 2 CHOICE:**
Multiple format options including textbook, eBook and eChapter rentals
- 3 CONVENIENCE:**
Anytime, anywhere access of eBooks or eChapters via mobile devices
- 4 SERVICE:**
Free eBook access while your text ships, and instant access to online homework products
- 5 STUDY TOOLS:**
Study tools* for your text, plus writing, research, career and job search resources
**availability varies*




Find your course materials and start saving at:
www.cengagebrain.com

Source Code: 14M-AA0107

ORACLE® 12c: SQL

Joan Casteel



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 to search by ISBN#, author, title, or keyword for materials in your areas of interest.

Important Notice: Media content referenced within the product description or the product text may not be available in the eBook version.

ORACLE® 12c: SQL

Joan Casteel

Vice President, General Manager, Science,
Math & Quantitative Business: Balraj Kalsi

Product Director: Joe Sabatino

Product Manager: Jason Guylar

Content Developer: Lori Bradshaw, S4 Carlisle

Senior Product Assistant: Brad Sullender

Senior Marketing Manager: Eric La Scola

Marketing Coordinator: William Guiliani

Art and Cover Direction, Production
Management, and Composition:

Lumina Datamatics, Inc.

Intellectual Property

Analyst: Christina Ciaramella

Project Manager: Kathryn Kucharek

Manufacturing Planner: Ron Montgomery

Cover Image: © Chris Clor/Blend Images/Corbis

© 2016, 2010 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**

Further permissions questions can be emailed to
permissionrequest@cengage.com

Oracle is a registered trademark, and Oracle 12c, SQL Developer, and SQL*Plus are trademarks or registered trademarks of Oracle Corporation and/or its affiliates.

Some of the product names and company names used in this book have been used for identification purposes only and may be trademarks or registered trademarks of their respective manufacturers and sellers.

Information pertaining to Northwest Airlines was used with their express permission. No part of it may be reproduced or used in any form without prior written permission from Cengage Learning.

Library of Congress Control Number: 2015942548

ISBN: 978-1-305-25103-8

Cengage Learning

20 Channel Center Street

Boston, MA 02210

USA

Cengage Learning is a leading provider of customized learning solutions with employees residing in nearly 40 different countries and sales in more than 125 countries around the world. Find your local representative at **www.cengage.com**

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

To learn more about Cengage Learning Solutions, visit **www.cengage.com**

Purchase any of our products at your local college store or at our
preferred online store **www.cengagebrain.com**

To Scott, a true teacher—one who never stops learning

BRIEF CONTENTS

Preface	xiv
Chapter 1 <i>Overview of Database Concepts</i>	1
Chapter 2 <i>Basic SQL SELECT Statements</i>	25
Chapter 3 <i>Table Creation and Management</i>	57
Chapter 4 <i>Constraints</i>	103
Chapter 5 <i>Data Manipulation and Transaction Control</i>	141
Chapter 6 <i>Additional Database Objects</i>	181
Chapter 7 <i>User Creation and Management</i>	225
Chapter 8 <i>Restricting Rows and Sorting Data</i>	255
Chapter 9 <i>Joining Data from Multiple Tables</i>	295
Chapter 10 <i>Selected Single-Row Functions</i>	347
Chapter 11 <i>Group Functions</i>	401
Chapter 12 <i>Subqueries and MERGE Statements</i>	449
Chapter 13 <i>Views</i>	495
Appendix A <i>Tables for the JustLee Books Database</i>	539
Appendix B <i>SQL*Plus and SQL Developer Overview</i>	547
Appendix C <i>Oracle Resources</i>	555
Appendix D <i>SQL*Loader</i>	557
Appendix E <i>SQL Tuning Topics</i>	561
Appendix F <i>SQL in Various Databases</i>	579
Glossary	583
Index	589

TABLE OF CONTENTS

Preface	xiv
Chapter 1 <i>Overview of Database Concepts</i>	1
Introduction	2
Database Terminology	2
Database Management System	3
Database Design	4
Entity-Relationship (E-R) Model	5
Database Normalization	6
Relating Tables in the Database	10
Structured Query Language (SQL)	12
Databases Used in This Textbook	13
Basic Assumptions	13
Tables in the JustLee Books Database	14
Topic Sequence	16
Software Used in This Textbook	16
Chapter Summary	17
Review Questions	17
Multiple Choice	18
Hands-On Assignments	21
Advanced Challenge	22
Case Study: <i>City Jail</i>	22
Chapter 2 <i>Basic SQL SELECT Statements</i>	25
Introduction	26
Creating the JustLee Books Database	27
SELECT Statement Syntax	30
Selecting All Data in a Table	31
Selecting One Column from a Table	33
Selecting Multiple Columns from a Table	34
Operations in the SELECT Statement	36
Using Column Aliases	36
Using Arithmetic Operations	39
NULL Values	40
Using DISTINCT and UNIQUE	42
Using Concatenation	44
Chapter Summary	49
Chapter 2 Syntax Summary	49

Review Questions	50
Multiple Choice	51
Hands-On Assignments	54
Advanced Challenge	55
Case Study: <i>City Jail</i>	55
Chapter 3 Table Creation and Management	57
Introduction	58
Table Design	59
Table Creation	63
Defining Columns	63
Viewing a List of Tables: USER_TABLES	65
Viewing Table Structures: DESCRIBE	66
Table Creation with Subqueries	70
CREATE TABLE ... AS Command	70
Modifying Existing Tables	72
ALTER TABLE ... ADD Command	73
ALTER TABLE ... MODIFY Command	74
ALTER TABLE ... DROP COLUMN Command	79
ALTER TABLE ... SET UNUSED/DROP UNUSED COLUMNS Command	80
Renaming a Table	83
Truncating a Table	85
Deleting a Table	86
Chapter Summary	91
Chapter 3 Syntax Summary	92
Review Questions	93
Multiple Choice	94
Hands-On Assignments	97
Advanced Challenge	97
Case Study: <i>City Jail</i>	98
Chapter 4 Constraints	103
Introduction	104
Creating Constraints	105
Creating Constraints at the Column Level	106
Creating Constraints at the Table Level	106
Using the PRIMARY KEY Constraint	107
Using the FOREIGN KEY Constraint	110
Using the UNIQUE Constraint	115
Using the CHECK Constraint	116
Using the NOT NULL Constraint	119
Including Constraints During Table Creation	121
Adding Multiple Constraints on a Single Column	125
Viewing Constraint Information	126

Disabling and Dropping Constraints	128
Using DISABLE/ENABLE	128
Dropping Constraints	129
Chapter Summary	131
Chapter 4 Syntax Summary	131
Review Questions	133
Multiple Choice	134
Hands-On Assignments	137
Advanced Challenge	139
Case Study: <i>City Jail</i>	139
Chapter 5 <i>Data Manipulation and Transaction Control</i>	141
Introduction	142
Inserting New Rows	143
Using the INSERT Command	143
Handling Virtual Columns	150
Handling Single Quotes in an INSERT Value	152
Inserting Data from an Existing Table	154
Modifying Existing Rows	156
Using the UPDATE Command	156
Using Substitution Variables	158
Deleting Rows	162
Using Transaction Control Statements	163
COMMIT and ROLLBACK Commands	164
SAVEPOINT Command	165
Using Table Locks	168
LOCK TABLE Command	168
SELECT ... FOR UPDATE Command	169
Chapter Summary	171
Chapter 5 Syntax Summary	172
Review Questions	173
Multiple Choice	173
Hands-On Assignments	176
Advanced Challenge	177
Case Study: <i>City Jail</i>	178
Chapter 6 <i>Additional Database Objects</i>	181
Introduction	182
Sequences	183
Creating a Sequence	184
Using Sequence Values	189
Setting and Altering Sequence Definitions	192
Altering Sequence Definitions	194
Removing a Sequence	196
Use Identity Columns Instead of Sequences for Primary Key Columns	197

Indexes	199
B-Tree Indexes	200
Bitmap Indexes	206
Function-Based Indexes	208
Index Organized Tables	209
Verifying an Index	210
Altering or Removing an Index	211
Synonyms	212
Deleting a Synonym	215
Chapter Summary	216
Chapter 6 Syntax Summary	217
Review Questions	218
Multiple Choice	219
Hands-On Assignments	222
Advanced Challenge	223
Case Study: <i>City Jail</i>	224
Chapter 7 <i>User Creation and Management</i>	225
Introduction	226
Data Security	227
Creating a User	228
Creating Usernames and Passwords	228
Assigning User Privileges	230
System Privileges	230
Granting System Privileges	231
Object Privileges	232
Granting Object Privileges	232
Managing Passwords	236
Using Roles	237
Creating and Assigning Roles	238
Using Predefined Roles	240
Using Default Roles	241
Enabling Roles After Login	242
Viewing Privilege Information	242
Removing Privileges and Users	244
Revoking Privileges and Roles	244
Dropping a Role	246
Dropping a User	246
Chapter Summary	247
Chapter 7 Syntax Summary	247
Review Questions	249
Multiple Choice	249
Hands-On Assignments	253
Advanced Challenge	253
Case Study: <i>City Jail</i>	254

Chapter 8	<i>Restricting Rows and Sorting Data</i>	255
Introduction		256
Where Clause Syntax		257
Rules for Character Strings		258
Rules for Dates		260
Comparison Operators		260
BETWEEN ... AND Operator		267
IN Operator		268
LIKE Operator		270
Logical Operators		274
Treatment of Null Values		277
ORDER by Clause Syntax		279
Secondary Sort		282
Sorting by SELECT Order		284
Chapter Summary		285
Chapter 8 Syntax Summary		286
Review Questions		288
Multiple Choice		288
Hands-On Assignments		292
Advanced Challenge		293
Case Study: <i>City Jail</i>		293
Chapter 9	<i>Joining Data from Multiple Tables</i>	295
Introduction		296
Cartesian Joins		297
Cartesian Join: Traditional Method		298
Cartesian Join: JOIN Method		300
Equality Joins		301
Equality Joins: Traditional Method		303
Equality Joins: JOIN Method		308
Non-Equality Joins		314
Non-Equality Joins: Traditional Method		315
Non-Equality Joins: JOIN Method		316
Self-Joins		317
Self-Joins: Traditional Method		318
Self-Joins: JOIN Method		319
Outer Joins		320
Outer Joins: Traditional Method		321
Outer Joins: JOIN Method		324
Set Operators		326
Chapter Summary		334
Chapter 9 Syntax Summary		335
Review Questions		337
Multiple Choice		338
Hands-On Assignments		344

Advanced Challenge	345
Case Study: <i>City Jail</i>	345
Chapter 10 Selected Single-Row Functions	347
Introduction	348
Case Conversion Functions	349
The LOWER Function	349
The UPPER Function	350
The INITCAP Function	351
Character Manipulation Functions	352
The SUBSTR Function	352
The INSTR Function	354
The LENGTH Function	356
The LPAD and RPAD Functions	357
The LTRIM and RTRIM Functions	358
The REPLACE Function	359
The TRANSLATE Function	360
The CONCAT Function	360
Number Functions	361
The ROUND Function	361
The TRUNC Function	362
The MOD Function	363
The ABS Function	364
The POWER Function	365
Date Functions	365
The MONTHS_BETWEEN Function	367
The ADD_MONTHS Function	367
The NEXT_DAY and LAST_DAY Functions	368
The TO_DATE Function	369
Rounding Date Values	371
Truncating Date Values	372
CURRENT_DATE Versus SYSDATE	373
Regular Expressions	374
Other Functions	377
The NVL Function	377
The NVL2 Function	380
The NULLIF Function	381
The TO_CHAR Function	383
The DECODE Function	385
The CASE Expression	387
The SOUNDEX Function	387
The TO_NUMBER Function	388
The DUAL Table	389
Chapter Summary	390
Chapter 10 Syntax Summary	390
Review Questions	394

Multiple Choice	395
Hands-On Assignments	399
Advanced Challenge	399
Case Study: <i>City Jail</i>	400
Chapter 11 <i>Group Functions</i>	401
Introduction	402
Group Functions	403
The SUM Function	404
The AVG Function	406
The COUNT Function	408
The MAX Function	411
The MIN Function	412
Grouping Data	413
Restricting Aggregated Output	417
Nesting Functions	421
Statistical Group Functions	422
The STDDEV Function	422
The VARIANCE Function	423
Enhanced Aggregation for Reporting	424
The GROUPING SETS Expression	427
The CUBE Extension	428
The ROLLUP Extension	431
Pattern Matching	437
Chapter Summary	439
Chapter 11 Syntax Summary	439
Review Questions	441
Multiple Choice	442
Hands-On Assignments	446
Advanced Challenge	447
Case Study: <i>City Jail</i>	447
Chapter 12 <i>Subqueries and MERGE Statements</i>	449
Introduction	450
Subqueries and Their Uses	451
Single-Row Subqueries	451
Single-Row Subquery in a WHERE Clause	451
Single-Row Subquery in a HAVING Clause	456
Single-Row Subquery in a SELECT Clause	457
Multiple-Row Subqueries	459
The IN Operator	460
The ALL and ANY Operators	461
Multiple-Row Subquery in a HAVING Clause	465
Multiple-Column Subqueries	467
Multiple-Column Subquery in a FROM Clause	467
Multiple-Column Subquery in a WHERE Clause	469

NULL Values	471
NVL in Subqueries	471
IS NULL in Subqueries	472
Correlated Subqueries	473
Nested Subqueries	475
Subquery Factoring Clause	477
DML Actions Using Subqueries	478
MERGE Statements	479
Chapter Summary	484
Chapter 12 Syntax Summary	484
Review Questions	486
Multiple Choice	486
Hands-On Assignments	492
Advanced Challenge	493
Case Study: <i>City Jail</i>	493
Chapter 13 Views	495
Introduction	496
Creating a View	498
Creating a Simple View	500
DML Operations on a Simple View	504
Creating a Complex View	508
DML Operations on a Complex View with an Arithmetic Expression	508
DML Operations on a Complex View Containing Data from Multiple Tables	513
DML Operations on a Complex View Containing Functions or Grouped Data	515
DML Operations on a Complex View Containing DISTINCT or ROWNUM	517
Summary Guidelines for DML Operations on a Complex View	519
Dropping a View	519
Creating an Inline View	520
CROSS and OUTER APPLY Methods for Joins	520
TOP-N Analysis	522
Creating a Materialized View	527
Chapter Summary	531
Chapter 13 Syntax Summary	532
Review Questions	533
Multiple Choice	533
Hands-On Assignments	537
Advanced Challenge	538
Case Study: <i>City Jail</i>	538
Appendix A Tables for the JustLee Books Database	539
CUSTOMERS Table	539
BOOKS Table	540
ORDERS Table	541
ORDERITEMS Table	542
AUTHOR Table	543
BOOKAUTHOR Table	544

PUBLISHER Table	545
PROMOTION Table	546
Appendix B <i>SQL*Plus and SQL Developer Overview</i>	547
Introduction	547
SQL*Plus	547
SQL Developer	551
Appendix C <i>Oracle Resources</i>	555
Oracle Academic Initiative (OAI)	555
Oracle Certification Program (OCP)	555
Oracle Technology Network (OTN)	555
International Oracle Users Group (IOUG)	556
Appendix D <i>SQL*Loader</i>	557
Introduction	557
Read a Fixed File Format	557
Read a Delimited File	559
Appendix E <i>SQL Tuning Topics</i>	561
Introduction	561
Tuning Concepts and Issues	561
Identifying Problem Areas in Coding	561
Processing and the Optimizer	563
The Explain Plan	565
Timing Feature	570
Selected SQL Tuning Guidelines and Examples	571
Avoiding Unnecessary Column Selection	572
Index Suppression	573
Concatenated Indexes	575
Subqueries	576
Optimizer Hints	577
Appendix F <i>SQL in Various Databases</i>	579
Introduction	579
Suppressing Duplicates	579
Locating a Value in a String	580
Displaying the Current Date	580
Specifying a Default Date Format	580
Replacing NULL Values in Text Data	581
Adding Time to Dates	581
Extracting Values from a String	581
Concatenating	582
Data Structures	582
Glossary	583
Index	589

PREFACE

The past few decades have seen a proliferation of organizations that rely heavily on information technology. These organizations store their data in databases, and many choose Oracle database management systems to access their data. The current Oracle database version, Oracle 12c, is a database management system that enables users to create, manipulate, and retrieve data. The purpose of this textbook is to introduce students to basic SQL commands for interacting with Oracle 12c databases in a business environment. In addition, concepts relating to objectives of the current Oracle certification exams have been incorporated for students wanting to pursue certification.

The Intended Audience

This textbook has been designed for students in technical two-year or four-year programs who need to learn how to interact with databases. Although having an understanding of database design is preferable, an introductory chapter has been included to review the basic concepts of E-R modeling and the normalization process.

Oracle Certification Program (OCP)

This textbook covers the objectives of SQL Fundamentals Exams available: 1Z0-051 for Oracle 11g and Exam 1Z0-061 for Oracle Database 12c. Most objectives for Exam 1Z0-047, Oracle Database SQL Expert, are also covered. Any of these exams serve as the first exam in the Oracle Application Development and Oracle Database Administrator Oracle Certified Associate level certification tracks. Information about registering for these exams, along with other reference material, is available at www.oracle.com.

The Approach

The concepts introduced in this textbook are discussed in the context of a hypothetical real-world business: an online book retailer named JustLee Books. The company's business operation and the database structure are introduced and analyzed, and as commands are introduced throughout the textbook, they're modeled with examples using the JustLee Books database. Using consistent examples of a hypothetical company helps you learn the syntax of commands and how to use them in a real-world environment. In addition, a script file that generates the database is available to give you hands-on practice in re-creating examples and practicing variations of SQL commands to enhance your understanding.

To explain what a database is and how it's created, this textbook initially focuses on creating tables and learning how to perform data manipulation operations. After you're familiar with the database structure, the focus then turns to querying a database.

In Chapters 8 through 13, you learn how to retrieve data from the database, using the many options of a SELECT statement, including row filtering, joins, functions, and subqueries.

To reinforce the material, each chapter includes a chapter summary and, when appropriate, a syntax guide for the commands covered in the chapter. In addition, each chapter includes review questions and hands-on activities that test your knowledge and challenge you to apply that knowledge to solving business problems. A running case study that builds throughout the textbook provides a second real-world setting—a city jail system—as another opportunity to work with databases.

New to This Edition

Topics added covering new Oracle 12c features include:

- Extended data types
- Invisible columns
- Default column values
 - Specifying a value for NULL inserts
 - Based on a sequence value
- Identity columns
- Left outer join enhancements
- CROSS APPLY join method
- Pattern matching for trend analyses
- Enhanced Top-N-Query features

In addition, coverage of subquery factoring or use of the WITH clause has been added to Chapter 12.

Overview of This Book

The examples, assignments, and cases in this book help you achieve the following objectives:

- Issue SQL commands that retrieve data based on criteria specified by the user.
- Use SQL commands to join tables and retrieve data from joined tables.
- Perform calculations based on data stored in the database.
- Use functions to manipulate and aggregate data.
- Use subqueries to retrieve data based on unknown conditions.
- Create, modify, and drop database tables.
- Manipulate data stored in database tables.
- Enforce business rules by using table constraints.
- Create users and assign the privileges users need to perform tasks.

The chapters' contents build in complexity while reinforcing previous ideas. Chapter 1 introduces basic database management concepts, including database design. Chapter 2 shows how to retrieve data from a table. Chapter 3 explains how to create new database tables. Chapter 4 addresses the use of constraints to enforce business rules and ensure the integrity of table data. Chapter 5 explains adding data to a table, modifying existing data, and deleting data. Chapter 6 shows how to use a sequence to generate numbers, create

indexes to speed up data retrieval, and create synonyms to provide aliases for tables. Chapter 7 steps you through creating user accounts and roles and shows how to grant (and revoke) privileges to these accounts and roles. Chapter 8 explains how to restrict rows retrieved from a table, based on a given condition. Chapter 9 shows how to link tables with common columns by using joins. Chapter 10 describes the single-row functions supported by Oracle 12c. Chapter 11 covers the usage of multiple-row functions to derive a single value for a group of rows and explains how to restrict groups of rows. **Chapter 12** covers using subqueries to retrieve rows based on an unknown condition already stored in the database. Chapter 13 explains using views to restrict access to data and reduce the complexity of certain types of queries.

The appendixes support and reinforce chapter materials. **Appendix A** contains printed versions of the initial table structure and data for the JustLee Books database used throughout this textbook. **Appendix B** introduces the operation of the SQL*Plus and SQL Developer client software tools. **Appendix C** lists Oracle resources for further study. **Appendix D** introduces the SQL*Loader utility for importing data. **Appendix E** introduces basic SQL statement tuning concepts. **Appendix F** identifies SQL differences in some popular databases.

Features

To enhance your learning experience, each chapter in this textbook includes the following elements:

- **Chapter objectives:** Each chapter begins with a list of the concepts to be mastered by the chapter's conclusion. This list gives you a quick overview of chapter contents and serves as a useful study aid.
- **Running case:** A sustained example, the business operation of JustLee Books, is the basis for introducing new commands and practicing the material covered in each chapter.
- **Methodology:** As new commands are introduced in each chapter, the command syntax is shown and then an example, using the JustLee Books database, illustrates using the command in the context of business operations. This methodology shows you not only *how* the command is used, but also *when* and *why* it's used. The script file used to create the database is available so that you can work through the examples in this textbook, engendering a hands-on environment in which you can reinforce your knowledge of chapter material.
- **Tip:** This feature, designated by the Tip icon, provides practical advice and sometimes explains how a concept applies in the workplace.
- **Note:** These explanations, designated by the Note icon, offer more information on performing operations with databases.
- **Database Preparation:** These notes, placed at the end-of-chapter introductions, tell you which script from the student data files you should run in preparation for chapter examples and activities.
- **Caution:** This warning, designated by the Caution icon, points out database operations that, if misused, could have devastating results.
- **Chapter summaries:** Each chapter's text is followed by a summary of chapter concepts. These summaries are a helpful recap of chapter contents.

- **Syntax summaries:** Beginning with Chapter 2, a Syntax Guide table is included after each chapter summary to recap the command syntax covered in the chapter.
- **Review questions:** End-of-chapter assessment begins with review questions that reinforce the main ideas introduced in each chapter. These questions ensure that you have mastered the concepts and understand the information covered in the chapter.
- **Multiple-choice questions:** Each chapter contains multiple-choice questions covering the material in the chapter. Oracle certification-type questions are included to prepare you for the type of questions you can expect on certification exams and measure your level of understanding.
- **Hands-on assignments:** Along with conceptual explanations and examples, each chapter includes hands-on assignments related to the chapter's contents. The purpose of these assignments is to give you practical experience. In most cases, the assignments are based on the JustLee Books database and build on the examples in the chapter.
- **Advanced challenge:** This section poses another problem about the JustLee Books database for you to solve and is larger in scope than the hands-on assignments.
- **Case studies:** At the end of each chapter is a major case study, designed to help you apply what you have learned to real-world situations. These cases give you the opportunity to synthesize and evaluate information independently, examine potential solutions, and make recommendations, much as you would in an actual business situation. These cases use a database based on a city jail system.

Supplemental Materials

The following supplemental materials are available when this book is used in a classroom setting. All teaching tools available with this book are provided to instructors on the Cengage Learning Web site at www.cengagebrain.com.

- **Electronic Instructor's Manual:** The Instructor's Manual accompanying this textbook includes the following items:
 - Additional instructional material to assist in class preparation, including suggestions for lecture topics
 - A sample syllabus
 - When applicable, information about potential problems that can occur in networked environments
- **Cengage Learning Testing Powered by Cognero is 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

- **PowerPoint presentations:** Microsoft® PowerPoint slides are included for each chapter. Instructors can use the slides in three ways: as teaching aids during classroom presentations, as printed handouts for classroom distribution, or as network-accessible resources for chapter review. Instructors can add their own slides for additional topics introduced to the class.
- **Data files:** The script files needed to create the JustLee Books and City Jail databases are available on the Cengage Learning Web site at www.cengagebrain.com.
- **Solution files:** Solutions to chapter examples, end-of-chapter review questions and multiple-choice questions, hands-on assignments, and case studies are available on the Cengage Learning Web site at www.cengagebrain.com. The solutions are password protected.
- **Figure files:** Figure files allow instructors to create their own presentations with figures from the textbook.

Acknowledgments

I feel fortunate that Cengage Learning pursued my authorship of this textbook and continues to support my efforts. I am one lucky person—I have two angels in heaven, my mother and grandmother, and one angel here on earth, Scott. Without them watching over me, I would not be able to tackle such challenges. I also want to thank my father, who always seems more excited than me every time I finish a book project.

However, this textbook is the result of an incredible effort by many people whom I wish I had the opportunity to thank personally. First, hats off to Lori Bradshaw and Wendy Langerud of S4Carlisle Publishing Services for managing the content for this project and keeping everyone on schedule. And a special thank you to Arul Joseph Raj of Lumina Datamatics LTD for his detailed efforts in preparing, reviewing, and formatting the content. There were many others involved in this project like product management, editing, marketing, and art development, and every effort was critical in making this book a reality. I truly appreciate all of these efforts.

In addition, I need to recognize the enormous contribution of colleagues and reviewers, who provided helpful suggestions and insight into the development of this textbook in earlier editions. And, finally, many thanks to the instructors who have adopted this text and have dedicated their efforts to assist others to learn these most important database skills.

READ THIS BEFORE YOU BEGIN

TO THE USER

Data Files

To work through the examples and complete the projects in this book, you need to load the data files created for this book. Your instructor will provide these data files, or you can download them from the Cengage Learning Web site at www.cengagebrain.com and then search for this book's title. The data files are designed to supply the same data shown in chapter examples, so you can have hands-on practice in re-creating the queries and their output. The tables in the database can be reset if you encounter problems, such as accidentally deleting data. Working through all examples is highly recommended to reinforce your learning.

Starting with Chapter 2, database script instructions are given at the beginning of the chapter, if applicable. These database script files are in the folder corresponding to the chapter (Chapter 5, Chapter 10, and so forth) on www.cengagebrain.com and have filenames such as `JLDB_Build_#.sql` (substituting the chapter number for the # symbol). If the computer in your school lab—or your own computer—has Oracle 12c installed, you can work through the chapter examples and complete the hands-on assignments and case projects. Many of the coding examples in this textbook can be completed successfully with previous versions of Oracle (Oracle 10g or Oracle 11g).

Connecting to Oracle

Most colleges will provide various options for accessing Oracle software to execute SQL statements required in this text, typically via the internet and/or using college computer labs. Contact your instructor first to identify the options available. If you choose to install Oracle 12c on your own computer, visit www.oracle.com and go to the database download area. The database software is freely available for learning purposes, however, it is very important that you view the documentation for the software prior to downloading and installing the software. Be sure to review the hardware and software requirements prior to attempting the install. You will be asked to setup a free Oracle Technology Network account before downloading. If you are new to Oracle, consult with your instructor prior to installing Oracle database software for guidance. This text was tested on Oracle Database 12c release 12.1.0.1.0.

When you install the Oracle software, you're prompted to change the password for certain default administrative user accounts. Make sure you record the accounts' names and passwords because you might need to log in to the database with one of these administrative accounts in later chapters. After you install Oracle, you're required to

enter a username and password to access the software. One default username created during installation is “scott.”

Also, note that you can't use your own computer to work through chapter examples and complete projects unless you have the data files. You can get these files from your instructor, or you can download them from the Cengage Learning Web site at www.cengagebrain.com and then search for this book's title. When you download the data files, they should be stored in a directory separate from any other files on your hard drive. You need to remember the path or folder containing these files because each script filename must be prefixed with its location before you run the script.

Visit Our Web Site

A supplemental chapter on formatting report output has been included in the book's online materials. Additional materials designed especially for this textbook might be available on the Cengage Learning Web site. Go to www.cengagebrain.com periodically and search this site for more details.

TO INSTRUCTORS

To complete examples and activities in this textbook, your students must have access to the data files included on the Instructor Resources GD (or downloaded from www.cengagebrain.com).

The data files consist of the JustLee Database folder and a folder for each chapter. Many chapters require running a script; if so, these instructions are given in a note at the beginning of the chapter. These scripts are in folders corresponding to the chapter (Chapter 5, Chapter 10, and so forth) and have filenames such as JLDB_Build_#.sql (substituting the chapter number for the # symbol). The initial database creation is done at the beginning of Chapter 2 to create the JustLee Books database. Students should run the scripts as instructed to have a copy of the tables stored in their schemas. You should instruct your students on how to access and copy data files to their own computers. The chapters and projects in this book were tested with Oracle 12c Standard Edition.

Cengage Learning Data Files

You are granted a license to copy data files to any computer or computer network used by people who have purchased this book.

CHAPTER 1

OVERVIEW OF DATABASE CONCEPTS

LEARNING OBJECTIVES

After completing this chapter, you will be able to do the following:

- Define database terms
- Identify the purpose of a database management system (DBMS)
- Explain database design by using entity-relationship models and normalization
- Explain the purpose of a Structured Query Language (SQL)
- Understand how this textbook's topics are sequenced and how the two sample databases are used
- Identify the software used in this textbook

INTRODUCTION

Imagine you're starting up an online book retail company. How will customer orders be recorded? Will customers be able to search for products by name or keywords? Will you be able to analyze sales information to track profits, determine product success, and target marketing efforts to customers? Analyzing thousands of orders could take days without using a database. A **database** simplifies these tasks because it's a storage structure that provides mechanisms for recording, manipulating, and retrieving data.

The database used throughout this textbook is based on the activities of a hypothetical business, an online bookseller named JustLee Books. The company sells books via the Internet to customers throughout the United States. When a new customer places an order, he or she provides data such as name, billing and shipping addresses, and items ordered. The company also uses a database for all books in inventory.

To access the data required for operating JustLee Books, management relies on a DBMS. A **database management system (DBMS)** is used to create and maintain the structure of a database, and then to enter, manipulate, and retrieve the data it stores. Creating an efficient database design is the key to using a database effectively to support an organization's business operations.

This chapter introduces basic database terminology and discusses the process of designing a database for JustLee Books.

DATABASE TERMINOLOGY

Whenever a customer opens an account with a company, certain data must be collected. In many cases, the customer completes an online form that asks for the customer's name, address, and so on, as shown in Figure 1-1.

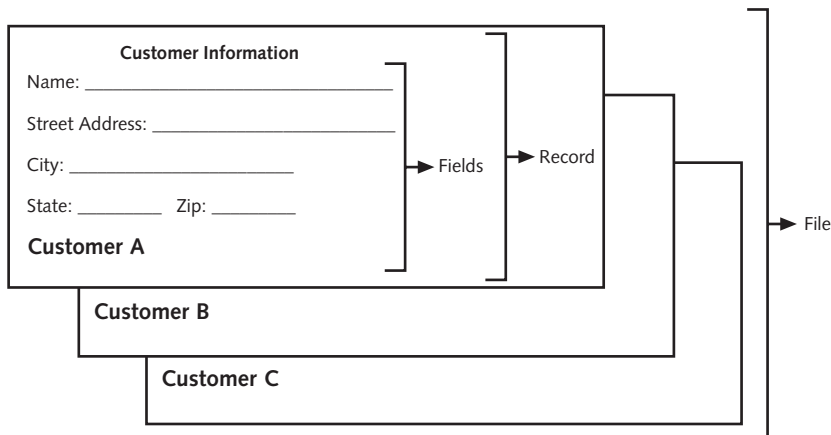


FIGURE 1-1 Collecting customer information

While collecting customer information, a series of characters is identified for each item. A **character** is the basic unit of data, and it can be a letter, number, or special symbol.

A group of related characters (for example, the characters that make up a customer's name) is called a **field**. A field represents one attribute or characteristic (the name, for instance) of the customer. A collection of fields about one customer (for example, name, address, city, state, and zip code) is called a **record**. A group of records about the same type of entity (such as customers or inventory items) is stored in a **file**. A collection of interrelated files—such as those relating to customers, their purchases, and their payments—is stored in a database.

These terms relate to the logical database design, but they are often used interchangeably with the terminology for the physical database design. When creating the physical database, a field is commonly referred to as a **column**, a record is called a **row**, and a file is known as a **table**. A table is quite similar to a spreadsheet, in that it contains columns and rows. Figure 1-2 shows a representation of these terms.

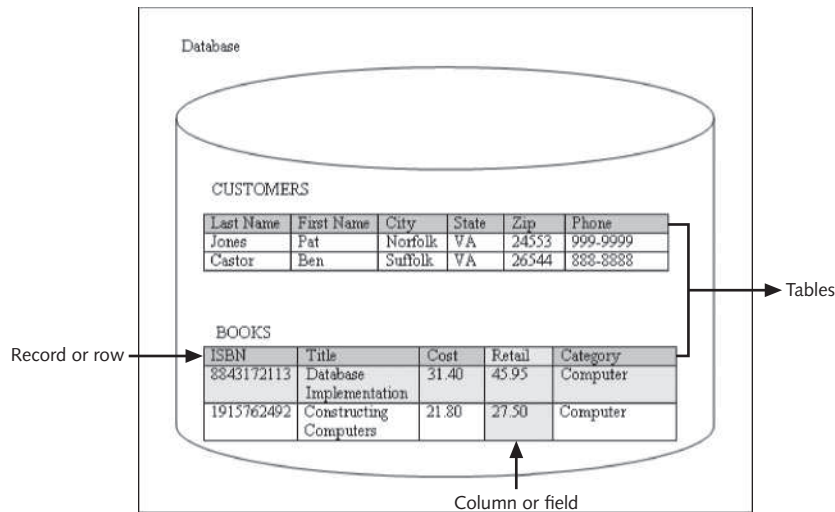


FIGURE 1-2 Database terminology

DATABASE MANAGEMENT SYSTEM

As mentioned earlier, a database is housed in a DBMS, which provides the functionality to create and work with a database. This functionality includes the following:

- *Data storage*: Manage the physical structure of the database.
- *Security*: Control user access and privileges.
- *Multiuser access*: Manage concurrent data access.
- *Backup*: Enable recovery options for database failures.
- *Data access language*: Provide a language that allows database access.
- *Data integrity*: Enable constraints or checks on data.
- *Data dictionary*: Maintain information about database structure.

DATABASE DESIGN

To determine the most appropriate structure of fields, records, and files in a database, developers go through a design process. The design and development of a system is accomplished through a process that's formally called the **Systems Development Life Cycle (SDLC)** and consists of the following steps:

1. *Systems investigation*: Understanding the problem
2. *Systems analysis*: Understanding the solution to the previously identified problem
3. *Systems design*: Defining the logical and physical components
4. *Systems implementation*: Creating the system
5. *Systems integration and testing*: Placing the system into operation for testing
6. *Systems deployment*: Placing the system into production
7. *Systems maintenance and review*: Evaluating the implemented system

Although the SDLC is a methodology designed for any type of system an organization needs, this chapter specifically addresses developing a DBMS. For the purposes of this discussion, assume the problem identified is the need to collect and maintain data about customers and their orders. The identified solution is to use a database to store all needed data. The discussion that follows presents the steps for designing the database.

NOTE

A variety of SDLC models have been developed to address different development environments. The steps presented here represent a traditional waterfall model. Other models, such as fountain and rapid prototyping, involve a different series of steps.

To design a database, the requirements of the database—inputs, processes, and outputs—must be identified first. Usually, the first question asked is, “What information, or output, must come from this database?” or “What questions should this database be able to answer?” By understanding the necessary output, the designer can then determine what information should be stored in the database. For example, if the organization wants to send birthday cards to its customers, the database must store each customer's birth date.

After the requirements of a database have been identified, an entity-relationship (E-R) model is usually drafted to better understand the data to be stored in the database. In an E-R model, an **entity** is any person, place, or thing with characteristics or attributes that will be included in the system. An **E-R model** is a diagram that identifies the entities (customers, books, orders, and such) in the database, and it shows how the entities are related to one another. It serves as the logical representation of the physical system to be built.

The next two sections explain the construction of an E-R model and the normalization process used to determine appropriate entities for a database.

NOTE

An E-R model is also called an entity-relationship diagram (ERD).

Entity-Relationship (E-R) Model

In an E-R model, an entity is usually represented as a square or rectangle. As shown in Figure 1-3, a line depicts how an entity's data relates to another entity. If the line connecting two entities is solid, the relationship between the entities is mandatory. However, if the relationship between two entities is optional, a dashed line is used.




FIGURE 1-3 E-R model notation examples

As shown in Figure 1-3, the following types of relationships can exist between two entities:

- *One-to-one*: In a one-to-one relationship, each occurrence of data in one entity is represented by only one occurrence of data in the other entity. For example, if each classroom is assigned to only one academic division, a one-to-one relationship is created between the classroom and division entities. This type of relationship is depicted in an E-R model as a simple straight line.
- *One-to-many*: In a one-to-many relationship, each occurrence of data in one entity can be represented by many occurrences of the data in the other entity. For example, a class has only one instructor, but an instructor might teach many classes. A one-to-many relationship is represented by a straight line with a crow's foot at the "many" end.
- *Many-to-many*: In a many-to-many relationship, data can have multiple occurrences in both entities. For example, a class can consist of more than one student, and a student can take more than one class. A straight line with a crow's foot at each end indicates a many-to-many relationship.

Figure 1-4 shows a simplified E-R model for the JustLee Books database used throughout this textbook. A more thorough E-R model would include a list of attributes for each entity.

NOTE

The notations in the sample E-R models in this chapter reflect only one way of diagramming entity relationships. If you're using a modeling software tool, you might encounter different notations to represent relationships. For example, Microsoft products typically represent the many side of a relationship with the infinity symbol (∞). In addition, some modeling tools automatically add the common fields or foreign key columns needed as relationships are defined.




FIGURE 1-4 An E-R model for JustLee Books

The following relationships are defined in the E-R model in Figure 1-4:

- Customers can place multiple orders, but each order can be placed by only one customer (one-to-many). The dashed line between Customers and Orders means a customer can exist in the database without having a current order stored in the ORDERS table. Therefore, this relationship is considered optional.
- An order can consist of more than one book, and a book can appear on more than one order (many-to-many).
- A book can have more than one author, and an author can write more than one book (many-to-many).
- A book can have only one publisher, but a publisher can publish more than one book (one-to-many).

Although some E-R modeling approaches are more complex, the simplified notations used in this chapter do point out the important relationships among entities, and using them helps designers identify potential problems in table layouts. After examining the E-R model in Figure 1-4, you should have noticed the two many-to-many relationships. Before creating the database, all many-to-many relationships must be reduced to a set of one-to-many relationships, as you learn in “Relating Tables in the Database” later in this chapter.

Identifying entities and relationships in the database design process is important because entities are usually represented as a table, and relationships can reveal whether additional tables are needed in the database. If the problem arising from the many-to-many relationship in the E-R model isn’t apparent to a designer at this point, it will become clear during the normalization process.

Database Normalization

Many people unfamiliar with database design principles often ask, “Why not just put all the data in one big table?” This single-table approach leads to problems of data redundancy (duplication) and data anomalies (data inconsistencies). For example, review the order data recorded in Table 1-1. The customer information is repeated for each order a customer places (redundancy). Also, the city data in the last row is different from the first two rows. Under these circumstances, it isn’t clear whether the last row actually represents a different

customer, whether the previous customer had an address change, or whether the city information is incorrect (data anomaly).

TABLE 1-1 Single-Table Approach Example

Last Name	First Name	City	State	Zip	Order Date	Order #
Jones	Pat	Norfolk	VA	24553	3/22/2009	45720
Jones	Pat	Norfolk	VA	24553	5/28/2009	48243
Jones	Pat	Suffolk	VA	26544	9/05/2009	51932

To avoid these data issues, database **normalization** is used to create a design that reduces or eliminates data redundancy and, therefore, avoids data anomalies. In general, normalization helps database designers determine which attributes, or fields, belong to each entity. In turn, this information helps determine which fields belong in each table. Normalization is a multistage process that enables designers to take the raw data to be collected about an entity and develop the data into a structured, normalized form that reduces the risks associated with data redundancy. Data redundancy poses a special problem in databases because storing the same data in different places can cause problems when updates or changes to data are required.

Most novices have difficulty understanding the impact of storing **unnormalized** data—data that hasn't been designed by using a normalization process. Here's an example. Suppose you work for a large company and submit a change-of-address form to the Human Resources (HR) Department. If all the data HR stores is normalized, a data entry clerk needs to update only the EMPLOYEES master table with your new address.

However, if the data is *not* stored in a normalized format, the data entry clerk likely needs to enter the change in each table containing your address—the EMPLOYEE RECORD table, the HEALTH INSURANCE table, the SICK LEAVE table, the ANNUAL TAX INFORMATION table, and so on—even though all this data is stored in the same database. As a result, if your mailing address is stored in several tables (or even duplicated in the same table) and the data entry clerk fails to make the change in one table, you might get a paycheck showing one address and, at the end of the year, have your W-2 form mailed to a different address! Storing data in a normalized format means only one update is required to reflect the new address, and it should always be the one that appears whenever your mailing address is needed.

A portion of the database for JustLee Books is used in this section to step through the normalization process—specifically, the books sold to customers. For each book, you need to store its International Standard Book Number (ISBN), title, publication date, wholesale cost, retail price, category (literature, self-help, and so forth), publisher name, contact person at the publisher for reordering the book (and telephone number), and author or authors' names.

Table 1-2 shows a sample of the data that must be maintained. For ease of illustration, the publishers' telephone numbers are eliminated, and the authors' names use just the first initial and last name. The first step in determining which data should be stored in each table is identifying a **primary key**, which is a field that identifies each record