



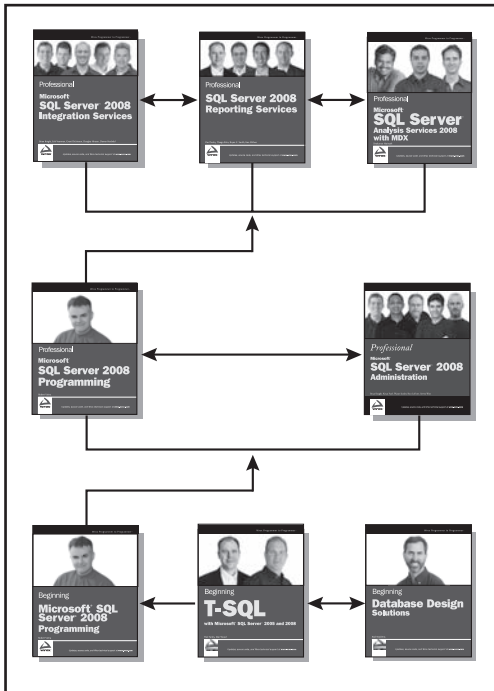
Beginning

Database Design Solutions

Rod Stephens



Beginning Database Design Solutions



Enhance Your Knowledge
Advance Your Career

Professional Microsoft SQL Server 2008 Integration Services

978-0-470-24795-2

This book shows developers how to master the 2008 release of SSIS, covering topics including data warehousing with SSIS, new methods of managing the SSIS platform, and improved techniques for ETL operations.

Professional SQL Server 2008 Reporting Services

978-0-470-24201-8

This book teaches solutions architects, designers, and developers how to use Microsoft's reporting platform to create reporting and business intelligence solutions.

Professional Microsoft SQL Server 2008 Analysis Services

978-0-470-24798-3

Professional Microsoft SQL Server 2008 Analysis Services shows readers how to build data warehouses and multidimensional databases, query databases, and use Analysis Services and other components of SQL Server to provide end-to-end solutions.

Professional Microsoft SQL Server 2008 Programming

978-0-470-25702-9

This updated new edition of Wrox's best-selling SQL Server book has been expanded to include coverage of SQL Server 2008's new datatypes, new indexing structures, manageability features, and advanced time-zone handling.

Professional Microsoft SQL Server 2008 Administration

978-0-470-24796-9

A how-to guide for experienced database administrators, this book is loaded with unique tips, tricks, and workarounds for handling the most difficult SQL Server administration issues. The authors discuss data capture, performance studio, Query Governor, and new techniques for monitoring and policy management.

Beginning Microsoft SQL Server 2008 Programming

978-0-470-25701-2

This comprehensive introduction to SQL Server covers the fundamentals and moves on to discuss how to create and change tables, manage keys, write scripts, work with stored procedures, and much more.

Beginning T-SQL with Microsoft SQL Server 2005 and 2008

978-0-470-25703-6

Beginning T-SQL with Microsoft SQL Server 2005 and 2008 provides a comprehensive introduction to the T-SQL programming language, with concrete examples showing how T-SQL works with both SQL Server 2005 and SQL Server 2008.

Beginning Database Design Solutions

978-0-470-38549-4

Beginning Database Design Solutions introduces IT professionals—both DBAs and database developers—to database design. It explains what databases are, their goals, and why proper design is necessary to achieve those goals. It tells how to decide what should be in a database to meet the application's requirements. It tells how to structure the database so it gives good performance while minimizing the chance for error.

Beginning Database Design Solutions

Introduction	xxiii
Part I: Introduction to Databases and Database Design	1
Chapter 1: Goals of Effective Database Design	3
Chapter 2: Database Types	23
Chapter 3: Relational Database Fundamentals	49
Part II: Database Design Process and Techniques	63
Chapter 4: Understanding User Needs	65
Chapter 5: Translating User Needs into Data Models	89
Chapter 6: Extracting Business Rules	121
Chapter 7: Normalizing Data	137
Chapter 8: Designing Databases to Support Software Applications	173
Chapter 9: Common Design Patterns	185
Chapter 10: Common Design Pitfalls	207
Part III: A Detailed Case Study	225
Chapter 11: User Needs and Requirements	227
Chapter 12: Building a Data Model	245
Chapter 13: Extracting Business Rules	263
Chapter 14: Normalization and Refinement	273
Part IV: Implementing Databases (with Examples in Access and MySQL)	285
Chapter 15: Microsoft Access	287
Chapter 16: MySQL	313
Part V: Advanced Topics	343
Chapter 17: Introduction to SQL	345
Chapter 18: Building Databases with SQL Scripts	369
Chapter 19: Database Maintenance	379
Chapter 20: Database Security	389
Appendix A: Exercise Solutions	403
Appendix B: Sample Database Designs	467
Glossary	487
Index	497

Beginning Database Design Solutions

Beginning Database Design Solutions

Rod Stephens



WILEY

Wiley Publishing, Inc.

Beginning Database Design Solutions

Published by
Wiley Publishing, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2009 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-470-38549-4

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Library of Congress Cataloging-in-Publication Data

Stephens, Rod, 1961-
Beginning database design solutions / Rod Stephens.
p. cm.
Includes index.
ISBN 978-0-470-38549-4 (978-0-470-38549-4)
1. Database design. 2. Databases. I. Title.
QA76.9.D26S97 2008
005.74 — dc22

2008037282

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Web site may provide or recommendations it may make. Further, readers should be aware that Internet Web sites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Trademarks: Wiley, the Wiley logo, Wrox, the Wrox logo, Wrox Programmer to Programmer, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

About the Author

Rod Stephens started out as a mathematician but, while studying at MIT, discovered the joys of computer algorithms and programming and he's been programming professionally ever since. During his career, he has worked on a wide variety of applications in such diverse fields as telephone switching, billing, repair dispatching, tax processing, wastewater treatment, concert ticket sales, cartography, and training for professional football players.

Rod is a Microsoft Visual Basic Most Valuable Professional (MVP), consultant and author. He has written 18 books that have been translated into half a dozen different languages, and more than 250 magazine articles, mostly about Visual Basic. Currently he is a regular contributor of C# and Visual Basic articles at DevX.com (www.devx.com).

Rod's popular *VB Helper* Web site www.vb-helper.com receives several million hits per month and contains thousands of pages of tips, tricks, and example code for Visual Basic programmers, as well as example code for this book.

Credits

Executive Editor

Robert Elliott

Development Editor

Sydney Jones

Technical Editor

Steve Hoberman

Production Editor

Angela Smith

Copy Editor

Kim Cofer

Editorial Manager

Mary Beth Wakefield

Production Manager

Tim Tate

Vice President and Executive Group Publisher

Richard Swadley

Vice President and Executive Publisher

Joseph B. Wikert

Project Coordinator, Cover

Lynsey Stanford

Proofreader

Publication Services, Inc.

Indexer

Jack Lewis

Acknowledgments

Thanks to Bob Elliott, Sydney Jones, Steve Hoberman, and all of the others whose hard work went into producing this book.

Special thanks to Sydney Jones for putting up with my sometimes overly generous interpretation of the Wrox guidelines.

Contents

Introduction

xxiii

Part I: Introduction to Databases and Database Design

Chapter 1: Goals of Effective Database Design **3**

Understanding the Importance of Design **4**

Information Containers **5**

Strengths and Weaknesses of Information Containers **7**

Desirable Database Features **8**

CRUD 9

Retrieval 9

Consistency 10

Validity 10

Easy Error Correction 11

Speed 11

Atomic Transactions 12

ACID 13

Persistence and Backups 14

Low Cost and Extensibility 16

Ease of Use 16

Portability 16

Security 17

Sharing 18

Ability to Perform Complex Calculations 19

Consequences of Good and Bad Design 19

Summary **21**

Exercises **22**

Chapter 2: Database Types **23**

Why Bother? **24**

Flat Files **24**

INI Files 26

Windows System Registry 26

Relational Databases **27**

Contents

Spreadsheets	30
Hierarchical Databases	31
XML	34
XML Basics	34
XML Structures	37
XML Summary	39
Network	40
Object	42
Object-Relational	43
Exotic	44
Document-Oriented	44
Deductive	44
Dimensional	44
Temporal	45
Summary	45
Exercises	47
Chapter 3: Relational Database Fundamentals	49
Relational Points of View	49
Table, Rows, and Columns	50
Relations, Attributes, and Tuples	52
Keys	52
Indexes	54
Constraints	55
Basic Constraints	55
Check Constraints	55
Primary Key Constraints	55
Unique Constraints	56
Foreign Key Constraints	56
Database Operations	57
Summary	59
Exercises	59
Part II: Database Design Process and Techniques	
Chapter 4: Understanding User Needs	65
Make a Plan	66
Bring a List of Questions	67
Functionality	67
Data Needs	67
Data Integrity	68

Security	68
Environment	69
Meet the Customers	69
Learn Who's Who	70
Pick the Customers' Brains	73
Walk a Mile in the User's Shoes	73
Study Current Operations	74
Brainstorm	75
Look to the Future	76
Understand the Customers' Reasoning	76
Learn What the Customers Really Need	77
Prioritize	78
Verify Your Understanding	79
Write the Requirements Document	80
Make Use Cases	81
Decide Feasibility	85
Summary	85
Exercises	86
Chapter 5: Translating User Needs into Data Models	89
What Are Data Models?	90
User Interface Models	92
Semantic Object Models	96
Classes and Objects	96
Cardinality	97
Identifiers	97
Putting It Together	97
Semantic Views	99
Class Types	100
Comments and Notes	106
Entity-Relationship Models	106
Entities, Attributes, and Identifiers	107
Relationships	107
Cardinality	108
Inheritance	109
Additional Conventions	112
Comments and Notes	113
Relational Models	113
Converting Semantic Object Models	113
Converting ER Diagrams	116
Summary	117
Exercises	118

Chapter 6: Extracting Business Rules	121
What Are Business Rules?	121
Identifying Key Business Rules	123
Extracting Key Business Rules	128
Multi-Tier Applications	129
Summary	133
Exercises	134
Chapter 7: Normalizing Data	137
What Is Normalization?	137
First Normal Form (1NF)	138
Second Normal Form (2NF)	146
Third Normal Form (3NF)	150
Stopping at Third Normal Form	154
Boyce-Codd Normal Form (BCNF)	154
Fourth Normal Form (4NF)	158
Fifth Normal Form (5NF)	162
Domain/Key Normal Form (DKNF)	165
Essential Redundancy	167
The Best Level of Normalization	168
Summary	169
Exercises	170
Chapter 8: Designing Databases to Support Software Applications	173
Plan Ahead	173
Document Everything	174
Consider Multi-Tier Architecture	175
Convert Domains into Tables	175
Keep Tables Focused	176
Use Three Kinds of Tables	176
Use Naming Conventions	178
Allow Some Redundant Data	179
Don't Squeeze in Everything	180
Summary	181
Exercises	182
Chapter 9: Common Design Patterns	185
Associations	185
Many-to-Many Associations	186
Multiple Many-to-Many Associations	186

Multiple-Object Associations	188
Repeated Attribute Associations	191
Reflexive Associations	193
Temporal Data	200
Effective Dates	200
Deleted Objects	202
Deciding What to Temporalize	203
Logging and Locking	203
Audit Trails	203
Turnkey Records	204
Summary	205
Exercises	205

Chapter 10: Common Design Pitfalls **207**

Lack of Preparation	207
Poor Documentation	208
Poor Naming Standards	208
Thinking Too Small	210
Not Planning for Change	210
Too Much Normalization	213
Insufficient Normalization	213
Insufficient Testing	214
Performance Anxiety	214
Mishmash Tables	215
Not Enforcing Constraints	217
Obsession with IDs	218
Not Defining Natural Keys	220
Summary	221
Exercises	222

Part III: A Detailed Case Study

Chapter 11: User Needs and Requirements **227**

Meet the Customers	227
Pick the Customers' Brains	229
Determining What the System Should Do	229
Determining How the Project Should Look	231
Determining What Data Is Needed for the User Interface	232
Determining Where the Data Should Come From	233
Determining How the Pieces of Data Are Related	233
Determining Performance Needs	235

Contents

Determining Security Needs	236
Determining Data Integrity Needs	236
Write Use Cases	238
Write the Requirements Document	242
Demand Feedback	243
Summary	244
Exercises	244
Chapter 12: Building a Data Model	245
Semantic Object Modeling	245
Building an Initial Semantic Object Model	245
Improving the Semantic Object Model	248
Entity-Relationship Modeling	250
Building an ER Diagram	250
Building a Combined ER Diagram	252
Improving the Entity-Relationship Diagram	254
Relational Modeling	256
Putting It All Together	259
Summary	260
Exercises	261
Chapter 13: Extracting Business Rules	263
Identifying Business Rules	263
Courses	264
CustomerCourses	266
Customers	267
Pets	267
Employees	267
Orders	267
OrderItems	268
InventoryItems	268
TimeEntries	268
Shifts	269
Persons	269
Phones	269
Vendors	269
Drawing a New Relational Model	270
Summary	270
Exercises	271

Chapter 14: Normalization and Refinement	273
Improving Flexibility	273
Verifying First Normal Form	275
Verifying Second Normal Form	278
Pets	279
TimeEntries	280
Verifying Third Normal Form	281
Summary	283
Exercises	284
Part IV: Implementing Databases	
	(with Examples in Access and MySQL)
Chapter 15: Microsoft Access	287
Understanding Access	287
Getting Started	288
Defining Relationships	293
Creating Field Constraints	296
Creating Table Constraints	298
Creating Queries	301
Query Design View	304
SQL View	306
Summary	308
Exercises	308
Chapter 16: MySQL	313
Installing MySQL	313
Using MySQL Command Line Client	314
Executing SQL Scripts	318
Using MySQL Query Browser	319
Executing Queries	320
Editing Data	320
Creating and Modifying Databases	321
Using Scripts	327
Getting Syntax Help	329
Using MySQL Workbench	330
Loading Scripts	330
Creating EER Diagrams	331

Contents

Editing Databases	333
Defining Triggers	334
Exporting Scripts	336
Summary	338
Exercises	338

Part V: Advanced Topics

Chapter 17: Introduction to SQL **345**

Background	345
Finding More Information	346
Standards	346
Basic Syntax	347
Command Overview	348
CREATE TABLE	350
CREATE INDEX	355
DROP	356
INSERT	357
SELECT	358
SELECT Clause	358
FROM Clause	359
WHERE Clause	362
GROUP BY Clause	362
ORDER BY Clause	363
UPDATE	364
DELETE	365
Summary	366
Exercises	366

Chapter 18: Building Databases with SQL Scripts **369**

Why Bother with Scripts?	369
Script Categories	370
Database Creation Scripts	370
Basic Initialization Scripts	370
Data Initialization Scripts	370
Cleanup Scripts	371
Saving Scripts	371
Ordering SQL Commands	371
Summary	376
Exercises	377

Chapter 19: Database Maintenance	379
Backups	379
Data Warehousing	382
Repairing the Database	383
Compacting the Database	383
Performance Tuning	383
Summary	386
Exercises	387
Chapter 20: Database Security	389
The Right Level of Security	389
Passwords	390
Single-Password Databases	390
Individual Passwords	390
Operating System Passwords	391
Good Passwords	391
Privileges	392
Initial Configuration and Privileges	398
Too Much Security	398
Physical Security	399
Summary	400
Exercises	401
Appendix A: Exercise Solutions	403
Chapter 1	403
Chapter 2	408
Chapter 3	411
Chapter 4	413
Chapter 5	417
Chapter 6	424
Chapter 7	430
Chapter 8	434
Chapter 9	436
Chapter 10	440
Chapter 11	442
Chapter 12	446
Chapter 13	449
Chapter 14	451
Chapter 15	454

Contents

Chapter 16	454
Chapter 17	454
Chapter 18	457
Chapter 19	459
Chapter 20	460
Appendix B: Sample Database Designs	467
Books	468
DVD and Movies	470
Music and CDs	471
Documents	472
Customer Orders	472
Employee Shifts and Timesheets	473
Employees, Projects, and Departments	473
Employee Skills and Qualifications	475
Identical Object Rental	476
Distinct Object Rental	476
Students, Courses, and Grades	478
Teams	479
Individual Sports	480
Vehicle Fleets	480
Contacts	483
Passengers	483
Recipes	485
Glossary	487
Index	497

Introduction

It has been estimated that more than 80 percent of all computer programming is database-related. This is certainly easy to believe. After all, a database can be a powerful tool for doing exactly what computer programs do best: store, manipulate, and display data.

Even many programs that seem at first glance to have little to do with traditional business-oriented data use databases to make processing easier. In fact, looking back on more than 20 years of software development experience, I'm hard pressed to think of a single non-trivial application that I've worked on that didn't use some kind of database.

Not only do databases play a role in many applications, but they also often play a critical role. If the data is not properly stored, it may become corrupted and the program will be unable to use it meaningfully. If the data is not properly organized, the program may be unable to find what it needs in a reasonable amount of time.

Unless the database stores its data safely and effectively, the application will be useless no matter how well-designed the rest of the system may be. The database is like the foundation of a building; without a strong foundation, even the best crafted building will fail, sometimes spectacularly (the Leaning Tower of Pisa notwithstanding).

With such a large majority of applications relying so heavily on databases, you would expect everyone involved with application development to have a solid, formal foundation in database design and construction. Everyone including database designers, application architects, programmers, database administrators, and project managers should ideally understand what makes a good database design. Even an application's key customers and users could benefit from understanding how databases work.

Sadly that is usually not the case. Many IT professionals have learned what they know about databases through rumor, trial-and-error, and painful experience. Over the years, some develop an intuitive feel for what makes a good database design but they may still not understand the reasons why a design is good or bad, and they may leave behind a trail of rickety, poorly constructed programs built on shaky database foundations.

This book provides the tools you need to design a database. It explains how to determine what should go in a database and how a database should be organized to ensure data integrity and a reasonable level of performance. It explains techniques for designing a database that is strong enough to store data safely and consistently, flexible enough to allow the application to retrieve the data it needs quickly and reliably, and adaptable enough to accommodate a realistic amount of change.

Introduction

With the ideas and techniques described in this book, you will be able to build a strong foundation for database applications.

Who This Book Is For

This book is intended for IT professionals and students who want to learn how to design, analyze, and understand databases. The material will benefit those who want a better high-level understanding of databases such as proposal managers, architects, project managers, and even customers. The material will also benefit those who will actually design, build, and work with databases such as database designers, database administrators, and programmers. In many projects, these roles overlap so the same person may be responsible for working on the proposal, managing part of the project, and designing and creating the database.

This book is aimed at IT professionals and students of all experience levels. It does not assume that you have any previous experience with databases or programs that use them. It doesn't even assume that you have experience with computers. All you really need is a willingness and desire to learn.

What This Book Covers

This book explains database design. It tells how to plan a database's structure so the database will be robust, resistant to errors, and flexible enough to accommodate a reasonable amount of future change. It explains how to discover database requirements, build data models to study data needs, and refine those models to improve the database's effectiveness.

The book solidifies these concepts by working through a detailed example that designs a realistic database. Later chapters explain how to actually build databases using two common database products: Access 2007 and MySQL.

The book finishes by describing some of the topics you need to understand to keep a database running effectively such as database maintenance and security.

What You Need to Use This Book

This book explains database design. It tells how to determine what should go in a database and how the database should be structured to give the best results.

This book does not focus on actually *creating* the database. The details of database construction are different for different database tools so, to remain as generally useful as possible, this book doesn't concentrate on any particular database system. You can apply the techniques described here equally to whatever database tool you use, whether it's Access, SQL Server, Oracle, MySQL, or some other database product.

Most database products include free editions that you can use for smaller projects. For example, SQL Server Express Edition, Oracle Express Edition, and MySQL Community Server are all free.

To remain database neutral, the book does not assume you are using a particular database so you don't need any particular software or hardware. To work through the Exercises, all you really need is a pencil and some paper. You are welcome to type solutions into your computer if you like but you may actually find working with pencil and paper easier than using a graphical design tool to draw pictures, at least until you are comfortable with database design and are ready to pick a computerized design tool.

Chapter 15, "Microsoft Access," explains how to build databases using the Microsoft Access 2007 database product. If you want to follow along with the examples in that chapter and work through the Exercises, you need to have Microsoft Access 2007 installed (although other versions of Access will also work with a few differences). You can use any operating system that will run Microsoft Access 2007.

Similarly Chapter 16, "MySQL," explains how to build databases using the MySQL Community Server database product. If you want to follow this chapter's examples and work through them, you will need to install MySQL Community Server. You can use any operating system that will run MySQL.

To experiment with the SQL database language described in Chapter 17, "Introduction to SQL," and Chapter 18, "Building Databases with SQL Scripts," you need any database product that supports SQL (that includes pretty much all relational databases) running on any operating system.

How This Book Is Structured

The chapters in this book are divided into five parts plus appendixes. The chapters in each part are described here. If you have previous experience with databases, you can use these descriptions to decide which chapters to skim and which to read in detail.

Part I: Introduction to Databases and Database Design

The chapters in this part of the book provide background that is necessary to understand the chapters that follow. You can skim some of this material if it is familiar to you but don't take it too lightly. If you understand the fundamental concepts underlying database design, it will be easier to understand the point behind important design concepts presented later.

Chapter 1, "Goals of Effective Database Design," explains the reasons why people and organizations use databases. It explains a database's purpose and conditions that it must satisfy to be useful. This chapter also describes the basic ACID (Atomicity, Consistency, Isolation, Durability) and CRUD (Create, Read, Update, Delete) features that any good database should have. It explains in high-level general terms what makes a good database and what makes a bad database.

Chapter 2, "Database Types," explains some of the different types of databases that you might decide to use. These include flat files, spreadsheets, hierarchical databases (XML), object databases, and relational databases. The relational database is one of the most powerful and most commonly used forms of database so it is the focus of this book, but it is important to realize that there are alternatives that may be more appropriate under certain circumstances. This chapter gives some tips on deciding which kind of database might be best for a particular project.

Chapter 3, "Relational Database Fundamentals," explains basic relational database concepts such as tables, rows, and columns. It explains the common usage of relational database terms in addition to

the more technical terms that are sometimes used by database theorists. It describes different kinds of constraints that databases use to guarantee that the data is stored safely and consistently.

Part II: Database Design Process and Techniques

The chapters in this part of the book discuss the main pieces of database design. They explain how to understand what should be in the database, develop an initial design, separate important pieces of the database to improve flexibility, and refine and tune the design to provide the most stable and useful design possible.

Chapter 4, “Understanding User Needs,” explains how to learn about the users’ needs and gather user requirements. It tells how to study the users’ current operations, existing databases (if any), and desired improvements. It describes common questions that you can ask to learn about users’ operations, desires, and needs, and how to build the results into requirements documents and specifications. This chapter explains what use cases are and tells how to use them and the requirements to guide database design and to measure success.

Chapter 5, “Translating User Needs into Data Models,” introduces data modeling. It explains how to translate the user’s conceptual model and the requirements into other more precise models that define the database design rigorously. This chapter describes several database modeling techniques including user-interface models, semantic object models, entity-relationship diagrams, and relational models.

Chapter 6, “Extracting Business Rules,” explains how a database can handle business rules. It explains what business rules are, how they differ from database structure requirements, and how you can identify business rules. This chapter explains the benefits of separating business rules from the database structure and tells how to achieve that separation.

Chapter 7, “Normalizing Data,” explains one of the biggest tools in database design: normalization. Normalization techniques allow you to restructure a database to increase its flexibility and make it more robust. This chapter explains the various forms of normalization, emphasizing the stages that are most common and important: first, second, and third normal forms (1NF, 2NF, and 3NF). It explains how each of these kinds of normalization helps prevent errors and tells why it is sometimes better to leave a database slightly less normalized to improve performance.

Chapter 8, “Designing Databases to Support Software Applications,” explains how databases fit into the larger context of application design and lifecycle. This chapter explains how later development depends on the underlying database design. It discusses multi-tier architectures that can help decouple the application and database design so there can be at least some changes to either without requiring changes to the other.

Chapter 9, “Common Design Patterns,” explains some common patterns that are useful in many applications. Some of these techniques include implementing various kinds of relationships among objects, storing hierarchical and network data, recording temporal data, and logging and locking.

Chapter 10, “Common Design Pitfalls,” explains some common design mistakes that occur in database development. It describes problems that can arise from insufficient planning, incorrect normalization, and obsession with ID fields and performance.

Part III: A Detailed Case Study

If you follow all of the examples and exercises in the earlier chapters, by this point you will have seen all of the major steps for producing a good database design. However, it's often useful to see all of the steps in a complicated process put together in a continuous sequence. The chapters in this part of the book walk through a detailed case study following all of the phases of database design for the fictitious Pampered Pet database.

Chapter 11, "User Needs and Requirements," walks through the steps required to analyze the users' problem, define requirements, and create use cases. It describes interviews with fictitious customers that are used to identify the application's needs and translate them into database requirements.

Chapter 12, "Building a Data Model," translates the requirements gathered in the previous chapter into a series of data models that precisely define the database's structure. This chapter builds user-interface models, entity-relationship diagrams, semantic object models, and relational models to refine the database's initial design. The final relational models match the structure of a relational database fairly closely so they are easy to implement.

Chapter 13, "Extracting Business Rules," identifies the business rules embedded in the relational model constructed in the previous chapter. It shows how to extract those rules in order to separate them logically from the database's structure. This makes the database more robust in the face of future changes to the business rules.

Chapter 14, "Normalization and Refinement," refines the relational model developed in the previous chapter by normalizing it. It walks through several versions of the database that are in different normal forms. It then selects the degree of normalization that provides a reasonable tradeoff between robust design and acceptable performance.

Part IV: Implementing Databases (with examples in Access and MySQL)

Though this book focuses on abstract database concepts that do not depend on a particular database product, it's also worth spending at least some time on more concrete implementation issues. The chapters in this part of the book describe some of those issues and explain how to build databases with two different database products: Access 2007 and MySQL.

Chapter 15, "Microsoft Access," explains how to build a database with Microsoft Access 2007. This chapter doesn't cover everything there is to know about Access, it just explains enough to get started and to use Access to build non-trivial databases. You can use other versions of Access to work through this chapter, although the locations of menus, buttons, and other Access features are different in different versions.

Chapter 16, "MySQL," explains how to build a database with MySQL. This chapter tells where to download a free version of MySQL. It explains how to use the MySQL Command Line Client as well as some useful graphical tools including MySQL Query Browser and MySQL Workbench.

Part V: Advanced Topics

Although this book does not assume you have previous database experience, that doesn't mean it cannot cover some more advanced subjects. The chapters in this part of the book explain some more sophisticated topics that are important but not central to database design.

Chapter 17, "Introduction to SQL," provides an introduction to SQL (Structured Query Language). It explains how to use SQL commands to add, insert, update, and delete data. By using SQL, you can help insulate a program from the idiosyncrasies of the particular database product that it uses to store data.

Chapter 18, "Building Databases with SQL Scripts," explains how to use SQL scripts to build a database. It explains the advantages of this technique, such as the ability to create scripts to initialize a database before performing tests. It also explains some of the restrictions on this method, such as the fact that the user must create and delete tables in specific orders to satisfy table relationships.

Chapter 19, "Database Maintenance," describes some of the database maintenance issues that are part of any database application. Though performing and restoring backups, compressing tables, rebuilding indexes, and populating data warehouses are strictly not database design tasks, they are essential to any working application.

Chapter 20, "Database Security," explains database security issues. It explains the kinds of security that some database products provide. It also explains some additional techniques that can enhance database security such as using database views to appropriately restrict the users' access to data.

Appendixes

The book's appendixes provide additional reference material to supplement the earlier chapters.

Appendix A, "Exercise Solutions," gives solutions to Exercises so you can check your progress as you work through the book.

Appendix B, "Sample Database Designs," includes the designs for a variety of common database situations. These designs store information about such topics as books, movies, documents, customer orders, employee timekeeping, rentals, students, teams, and vehicle fleets.

The Glossary provides definitions for useful database and software development terms. The Glossary includes terms defined and used in this book in addition to other useful terms that you may encounter while reading other database material. This appendix can be a useful reference when you encounter an unfamiliar term on the Web or in database articles.

How to Use This Book

Because this book is aimed at readers of all experience levels, you may find some of the material familiar if you have previous experience with databases. In that case, you may want to skim chapters covering material that you already thoroughly understand.

If you are familiar with relational databases, you may want to skim Chapter 1, "Goals of Effective Database Design," Chapter 2, "Database Types," and Chapter 3, "Relational Database Fundamentals."