

T-SQL Fundamentals

Third Edition



Professional



Itzik Ben-Gan
 SolidQ

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



T-SQL Fundamentals, Third Edition

Itzik Ben-Gan

PUBLISHED BY
Microsoft Press
A division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2016 by Itzik Ben-Gan

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2015955815
ISBN: 978-1-5093-0200-0

Printed and bound in the United States of America.

First Printing

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Support at mspinput@microsoft.com. Please tell us what you think of this book at <http://aka.ms/tellpress>.

This book is provided “as-is” and expresses the author’s views and opinions. The views, opinions and information expressed in this book, including URL and other Internet website references, may change without notice.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

Microsoft and the trademarks listed at <http://www.microsoft.com> on the “Trademarks” webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

Acquisitions Editor: Devon Musgrave

Developmental Editor: Carol Dillingham

Project Editor: Carol Dillingham

Editorial Production: Christian Holdener; S4Carlisle Publishing Services

Technical Reviewer: Bob Beauchemin; Technical Review services provided by Content Master, a member of CM Group, Ltd.

Copyeditor: Roger Leblanc

Indexer: Maureen Johnson, MoJo’s Indexing Services

Cover: Twist Creative • Seattle and Joel Panchot

*To Dato,
To live in hearts we leave behind,
Is not to die.*

—THOMAS CAMPBELL

This page intentionally left blank

Contents at a glance

	<i>Introduction</i>	<i>xix</i>
CHAPTER 1	Background to T-SQL querying and programming	1
CHAPTER 2	Single-table queries	27
CHAPTER 3	Joins	103
CHAPTER 4	Subqueries	133
CHAPTER 5	Table expressions	161
CHAPTER 6	Set operators	193
CHAPTER 7	Beyond the fundamentals of querying	213
CHAPTER 8	Data modification	249
CHAPTER 9	Temporal tables	297
CHAPTER 10	Transactions and concurrency	319
CHAPTER 11	Programmable objects	361
	<i>Appendix: Getting started</i>	<i>395</i>
	<i>Index</i>	<i>415</i>

This page intentionally left blank

Contents

Introduction	xix
Chapter 1 Background to T-SQL querying and programming	1
Theoretical background	1
SQL	2
Set theory	3
Predicate logic	4
The relational model	4
Types of database systems	10
SQL Server architecture	12
The ABCs of Microsoft RDBMS flavors	12
SQL Server instances	14
Databases	15
Schemas and objects	18
Creating tables and defining data integrity	19
Creating tables	20
Defining data integrity	22
Conclusion	25
Chapter 2 Single-table queries	27
Elements of the <i>SELECT</i> statement	27
The <i>FROM</i> clause	29
The <i>WHERE</i> clause	31
The <i>GROUP BY</i> clause	32
The <i>HAVING</i> clause	35
The <i>SELECT</i> clause	36

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can improve our books and learning resources for you. To participate in a brief survey, please visit:

<http://aka.ms/tellpress>

The <i>ORDER BY</i> clause	41
The <i>TOP</i> and <i>OFFSET-FETCH</i> filters	43
A quick look at window functions	47
Predicates and operators	49
<i>CASE</i> expressions	52
<i>NULLs</i>	54
All-at-once operations	58
Working with character data	60
Data types	60
Collation	61
Operators and functions	63
The <i>LIKE</i> predicate	71
Working with date and time data	73
Date and time data types	74
Literals	74
Working with date and time separately	78
Filtering date ranges	79
Date and time functions	80
Querying metadata	89
Catalog views	90
Information schema views	91
System stored procedures and functions	91
Conclusion	93
Exercises	93
Exercise 1	93
Exercise 2	93
Exercise 3	94
Exercise 4	94
Exercise 5	95
Exercise 6	95
Exercise 7	95
Exercise 8	96
Exercise 9	96
Exercise 10	97

Solutions	97
Exercise 1	98
Exercise 2	98
Exercise 3	99
Exercise 4	99
Exercise 5	99
Exercise 6	100
Exercise 7	101
Exercise 8	102
Exercise 9	102
Exercise 10	102

Chapter 3 Joins 103

Cross joins	103
ISO/ANSI SQL-92 syntax	104
ISO/ANSI SQL-89 syntax	105
Self cross joins	105
Producing tables of numbers	106
Inner joins	107
ISO/ANSI SQL-92 syntax	107
ISO/ANSI SQL-89 syntax	108
Inner join safety	109
More join examples	109
Composite joins	110
Non-equi joins	111
Multi-join queries	113
Outer joins	113
Fundamentals of outer joins	114
Beyond the fundamentals of outer joins	116
Conclusion	123
Exercises	123
Exercise 1-1	123
Exercise 1-2 (optional, advanced)	124
Exercise 2	125

Exercise 3	125
Exercise 4	126
Exercise 5	126
Exercise 6	127
Exercise 7 (optional, advanced)	127
Exercise 8 (optional, advanced)	128
Exercise 9 (optional, advanced)	128
Solutions	129
Exercise 1-1	129
Exercise 1-2	129
Exercise 2	129
Exercise 3	130
Exercise 4	130
Exercise 5	131
Exercise 6	131
Exercise 7	131
Exercise 8	132
Exercise 9	132

Chapter 4 Subqueries 133

Self-contained subqueries	133
Self-contained scalar subquery examples	134
Self-contained multivalued subquery examples	136
Correlated subqueries	139
The <i>EXISTS</i> predicate	141
Beyond the fundamentals of subqueries	143
Returning previous or next values	143
Using running aggregates	144
Dealing with misbehaving subqueries	145
Conclusion	149
Exercises	150
Exercise 1	150
Exercise 2 (optional, advanced)	150
Exercise 3	151

Exercise 4	151
Exercise 5	152
Exercise 6	152
Exercise 7 (optional, advanced)	153
Exercise 8 (optional, advanced)	153
Exercise 9	154
Exercise 10 (optional, advanced)	154
Solutions	154
Exercise 1	154
Exercise 2	155
Exercise 3	155
Exercise 4	155
Exercise 5	156
Exercise 6	156
Exercise 7	156
Exercise 8	157
Exercise 9	157
Exercise 10	158

Chapter 5 Table expressions 161

Derived tables	161
Assigning column aliases	163
Using arguments	164
Nesting	165
Multiple references	166
Common table expressions	167
Assigning column aliases in CTEs	167
Using arguments in CTEs	168
Defining multiple CTEs	168
Multiple references in CTEs	169
Recursive CTEs	169
Views	171
Views and the <i>ORDER BY</i> clause	172
View options	175

Inline table-valued functions	178
The <i>APPLY</i> operator	180
Conclusion	183
Exercises	183
Exercise 1	183
Exercise 2-1	184
Exercise 2-2	184
Exercise 3-1	184
Exercise 3-2	185
Exercise 4 (optional, advanced)	185
Exercise 5-1	186
Exercise 5-2 (optional, advanced)	187
Exercise 6-1	187
Exercise 6-2	188
Solutions	188
Exercise 1	188
Exercise 2-1	189
Exercise 2-2	189
Exercise 3-1	189
Exercise 3-2	189
Exercise 4	190
Exercise 5-1	190
Exercise 5-2	191
Exercise 6-1	191
Exercise 6-2	192

Chapter 6 Set operators 193

The <i>UNION</i> operator	194
The <i>UNION ALL</i> operator	194
The <i>UNION (DISTINCT)</i> operator	195
The <i>INTERSECT</i> Operator	196
The <i>INTERSECT (DISTINCT)</i> operator	196
The <i>INTERSECT ALL</i> operator	197
The <i>EXCEPT</i> operator	199

The <i>EXCEPT (DISTINCT)</i> operator.....	199
The <i>EXCEPT ALL</i> operator.....	200
Precedence.....	201
Circumventing unsupported logical phases.....	203
Conclusion.....	204
Exercises.....	204
Exercise 1.....	204
Exercise 2.....	205
Exercise 3.....	205
Exercise 4.....	206
Exercise 5.....	207
Exercise 6 (optional, advanced).....	207
Solutions.....	208
Exercise 1.....	208
Exercise 2.....	208
Exercise 3.....	209
Exercise 4.....	210
Exercise 5.....	210
Exercise 6.....	211

Chapter 7 Beyond the fundamentals of querying 213

Window functions.....	213
Ranking window functions.....	216
Offset window functions.....	219
Aggregate window functions.....	221
Pivoting data.....	224
Pivoting with a grouped query.....	226
Pivoting with the <i>PIVOT</i> operator.....	227
Unpivoting data.....	230
Unpivoting with the <i>APPLY</i> operator.....	231
Unpivoting with the <i>UNPIVOT</i> operator.....	233
Grouping sets.....	234
The <i>GROUPING SETS</i> subclause.....	236
The <i>CUBE</i> subclause.....	236

The <i>ROLLUP</i> subclause	236
The <i>GROUPING</i> and <i>GROUPING_ID</i> functions	238
Conclusion	240
Exercises	241
Exercise 1	241
Exercise 2	241
Exercise 3	242
Exercise 4	242
Exercise 5	243
Exercise 6	244
Solutions	245
Exercise 1	245
Exercise 2	245
Exercise 3	245
Exercise 4	246
Exercise 5	247
Exercise 6	247

Chapter 8 Data modification 249

Inserting data	249
The <i>INSERT VALUES</i> statement	249
The <i>INSERT SELECT</i> statement	251
The <i>INSERT EXEC</i> statement	251
The <i>SELECT INTO</i> statement	252
The <i>BULK INSERT</i> statement	253
The identity property and the sequence object	254
Deleting data	262
The <i>DELETE</i> statement	263
The <i>TRUNCATE</i> statement	264
<i>DELETE</i> based on a join	265
Updating data	266
The <i>UPDATE</i> statement	267
<i>UPDATE</i> based on a join	268
Assignment <i>UPDATE</i>	270

Merging data	271
Modifying data through table expressions	276
Modifications with <i>TOP</i> and <i>OFFSET-FETCH</i>	278
The <i>OUTPUT</i> clause	280
<i>INSERT</i> with <i>OUTPUT</i>	280
<i>DELETE</i> with <i>OUTPUT</i>	282
<i>UPDATE</i> with <i>OUTPUT</i>	283
<i>MERGE</i> with <i>OUTPUT</i>	284
Nested DML	285
Conclusion	287
Exercises	287
Exercise 1	287
Exercise 2	288
Exercise 3	288
Exercise 4	288
Exercise 5	290
Exercise 6	290
Solutions	291
Exercise 1	291
Exercise 2	293
Exercise 3	293
Exercise 4	294
Exercise 5	294
Exercise 6	295

Chapter 9 Temporal tables 297

Creating tables	298
Modifying data	301
Querying data	304
Conclusion	311
Exercises	311
Exercise 1	311
Exercise 2	311

Exercise 3	312
Exercise 4	313
Solutions.....	313
Exercise 1.....	313
Exercise 2	314
Exercise 3	316
Exercise 4	317

Chapter 10 Transactions and concurrency 319

Transactions.....	319
Locks and blocking	322
Locks	322
Troubleshooting blocking.....	325
Isolation levels	332
The <i>READ UNCOMMITTED</i> isolation level.....	333
The <i>READ COMMITTED</i> isolation level.....	334
The <i>REPEATABLE READ</i> isolation level.....	335
The <i>SERIALIZABLE</i> isolation level	337
Isolation levels based on row versioning	339
Summary of isolation levels	345
Deadlocks.....	345
Conclusion	348
Exercises	348
Exercise 1-1	349
Exercise 1-2.....	349
Exercise 1-3.....	349
Exercise 1-4	349
Exercise 1-5.....	350
Exercise 1-6	350
Exercise 2-1.....	350
Exercise 2-2	351
Exercise 2-3	352
Exercise 2-4	353
Exercise 2-5	354

Exercise 2-6	356
Exercise 3-1.....	357
Exercise 3-2	358
Exercise 3-3	358
Exercise 3-4	358
Exercise 3-5	358
Exercise 3-6	358
Exercise 3-7	359

Chapter 11 Programmable objects 361

Variables	361
Batches	363
A batch as a unit of parsing	364
Batches and variables	365
Statements that cannot be combined in the same batch	365
A batch as a unit of resolution	366
The <i>GO n</i> option	366
Flow elements	367
The <i>IF...ELSE</i> flow element	367
The <i>WHILE</i> flow element	368
Cursors	370
Temporary tables	374
Local temporary tables	374
Global temporary tables	376
Table variables	377
Table types	378
Dynamic SQL	379
The <i>EXEC</i> command	380
The <i>sp_executesql</i> stored procedure	380
Using <i>PIVOT</i> with Dynamic SQL	381
Routines	383
User-defined functions	383
Stored procedures	385
Triggers	387

Error handling	390
Conclusion	394
<i>Appendix: Getting started</i>	395
<i>Index</i>	415

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can improve our books and learning resources for you. To participate in a brief survey, please visit:

<http://aka.ms/tellpress>

Introduction

This book walks you through your first steps in T-SQL (also known as *Transact-SQL*), which is the Microsoft SQL Server dialect of the ISO and ANSI standards for SQL. You'll learn the theory behind T-SQL querying and programming and how to develop T-SQL code to query and modify data, and you'll get an overview of programmable objects.

Although this book is intended for beginners, it's not merely a set of procedures for readers to follow. It goes beyond the syntactical elements of T-SQL and explains the logic behind the language and its elements.

Occasionally, the book covers subjects that might be considered advanced for readers who are new to T-SQL; therefore, you should consider those sections to be optional reading. The text will indicate when a section is considered more advanced and is provided as optional reading. If you feel comfortable with the material discussed in the book up to that point, you might want to tackle these more advanced subjects; otherwise, feel free to skip those sections and return to them after you gain more experience.

Many aspects of SQL are unique to the language and very different from other programming languages. This book helps you adopt the right state of mind and gain a true understanding of the language elements. You learn how to think in relational terms and follow good SQL programming practices.

The book is not version specific; it does, however, cover language elements that were introduced in recent versions of SQL Server, including SQL Server 2016. When I discuss language elements that were introduced recently, I specify the version in which they were added.

Besides being available as a box product, SQL Server is also available as a cloud-based service called Microsoft Azure SQL Database, or in short, just SQL Database. The code samples in this book were tested against both a box SQL Server product and Azure SQL Database. The book's companion content (available at <http://aka.ms/T-SQLFund3e/downloads>) provides information about compatibility issues between the flavors.

To complement the learning experience, the book provides exercises you can use to practice what you learn. The book occasionally provides optional exercises that are more advanced. Those exercises are intended for readers who feel comfortable with the material and want to challenge themselves with more difficult problems. The optional exercises for advanced readers are labeled as such.

Who should read this book

This book is intended for T-SQL developers, database administrators (DBAs), business intelligence (BI) practitioners, data scientists, report writers, analysts, architects, and SQL Server power users who just started working with SQL Server and who need to write queries and develop code using Transact-SQL.

Assumptions

To get the most out of this book, you should have working experience with Microsoft Windows and with applications based on Windows. You should also be familiar with basic concepts of relational database management systems.

This book might not be for you if...

Not every book is aimed at every possible audience. This book covers fundamentals. It's mainly aimed at T-SQL practitioners with little or no experience. This book might not be for you if you're an advanced T-SQL practitioner with many years of experience. With that said, several readers of the previous editions of this book have mentioned that—even though they already had years of experience—they still found the book useful for filling gaps in their knowledge.

Organization of this book

This book starts with a theoretical background to T-SQL querying and programming in Chapter 1, laying the foundations for the rest of the book, and basic coverage of creating tables and defining data integrity. The book moves on to various aspects of querying and modifying data in Chapters 2 through 9, and then moves to a discussion of concurrency and transactions in Chapter 10. Finally, it provides an overview of programmable objects in Chapter 11.

Here's a list of the chapters along with a short description of the content in each chapter:

- Chapter 1, "Background to T-SQL querying and programming," provides the theoretical background for SQL, set theory, and predicate logic. It examines the relational model, describes SQL Server's architecture, and explains how to create tables and define data integrity.

- Chapter 2, “Single-table queries,” covers various aspects of querying a single table by using the *SELECT* statement.
- Chapter 3, “Joins,” covers querying multiple tables by using joins, including cross joins, inner joins, and outer joins.
- Chapter 4, “Subqueries,” covers queries within queries, otherwise known as *subqueries*.
- Chapter 5, “Table expressions,” covers derived tables, Common Table Expressions (CTEs), views, inline table-valued functions, and the *APPLY* operator.
- Chapter 6, “Set operators,” covers the set operators *UNION*, *INTERSECT*, and *EXCEPT*.
- Chapter 7, “Beyond the fundamentals of querying,” covers window functions, pivoting, unpivoting, and working with grouping sets.
- Chapter 8, “Data modification,” covers inserting, updating, deleting, and merging data.
- Chapter 9, “Temporal tables,” covers system-versioned temporal tables.
- Chapter 10, “Transactions and concurrency,” covers concurrency of user connections that work with the same data simultaneously; it covers transactions, locks, blocking, isolation levels, and deadlocks.
- Chapter 11, “Programmable objects,” provides an overview of the T-SQL programming capabilities in SQL Server.
- The book also provides an appendix, “Getting started,” to help you set up your environment, download the book’s source code, install the TSQLV4 sample database, start writing code against SQL Server, and learn how to get help by working with SQL Server Books Online.

System requirements

The Appendix, “Getting started,” explains which editions of SQL Server 2016 you can use to work with the code samples included with this book. Each edition of SQL Server might have different hardware and software requirements, and those requirements are well documented in SQL Server Books Online under “Hardware and Software Requirements for Installing SQL Server 2016” at the following URL: <https://msdn.microsoft.com/en-us/library/ms143506.aspx>. The Appendix also explains how to work with SQL Server Books Online.

If you’re connecting to Azure SQL Database, hardware and server software are handled by Microsoft, so those requirements are irrelevant in this case.

You will need SQL Server Management Studio (SSMS) to run the code samples against both SQL Server 2016 and Azure SQL Database. You can download SSMS and find information about the supported operating systems at the following URL: <https://msdn.microsoft.com/en-us/library/mt238290.aspx>.

Installing and using the source code

Most of the chapters in this book include exercises that let you interactively try out new material learned in the main text. All source code, including exercises and solutions, can be downloaded from the following webpage:

<http://aka.ms/T-SQLFund3e/downloads>

Follow the instructions to download the TSQLFundamentalsYYYYMMDD.zip file, where YYYYMMDD reflects the last update date of the source code.

Refer to the Appendix, “Getting started,” for details about the source code.

Acknowledgments

A number of people contributed to making this book a reality, either directly or indirectly, and deserve thanks and recognition. It’s certainly possible I omitted some names unintentionally, and I apologize for this ahead of time.

To Lilach: You’re the one who makes me want to be good at what I do. Besides being my inspiration in life, you always take an active role in my books, helping to review the text for the first time.

To my mom, Mila, and to my siblings, Mickey and Ina: Thank you for the constant support and for accepting the fact that I’m away. To my dad, Gabi, who loved puzzles, logic, and numbers: I attribute my affinity to SQL to you; we all miss you a lot.

To the technical reviewer of the book, Bob Beauchemin: You’ve been around in the SQL Server community for many years. I was always impressed by your extensive knowledge and was happy that you agreed to work with me on this book.

To Steve Kass, Dejan Sarka, Gianluca Hotz, and Herbert Albert: Thanks for your valuable advice during the planning phases of the book. I had to make some hard decisions in terms of what to include and what not to include in the book, and your advice was very helpful.

Many thanks to the book’s editors. To Devon Musgrave, who played the acquisitions editor role: You are the one who made this book a reality and handled all the initial

stages. I realize that this book is likely one of many you were responsible for, and I'd like to thank you for dedicating the time and effort that you did. To Carol Dillingham, the book's developmental editor and project editor, many thanks for your excellent handling; I always enjoy working with you on my books. Also thanks to Roger LeBlanc for his fine copy edits, and to Christian Holdener for his project management.

To SolidQ, my company for over a decade: It's gratifying to be part of such a great company that evolved into what it is today. The members of this company are much more than colleagues to me; they are partners, friends, and family. To Fernando G. Guerrero, Antonio Soto, and Douglas McDowell: thanks for leading the company. To my many colleagues: It's an honor to be part of this amazingly talented group.

To members of the Microsoft SQL Server development team, past and present: Tobias Ternstrom, Lubor Kollar, Umachandar Jayachandran (UC), Boris Baryshnikov, Conor Cunningham, Kevin Farlee, Josde Bruijn, Marc Friedman, Drazen Sumic, Borko Novakovic, Milan Stojic, Milan Ruzic, Jovan Popovic, Lindsey Allen, Craig Freedman, Campbell Fraser, Eric Hanson, Mark Souza, Dave Campbell, César Galindo-Legaria, Pedro Lopes, and I'm sure many others. Thanks for creating such a great product, and thanks for all the time you spent meeting with me and responding to my emails, addressing my questions, and answering my requests for clarification.

To members of the SQL Server Pro editorial team, Tim Ford and Debra Donston-Miller: I've been writing for the magazine for almost two decades, and I'm grateful for the opportunity to share my knowledge with the magazine's readers.

To Data Platform MVPs, past and present: Paul White, Alejandro Mesa, Erland Sommarskog, Aaron Bertrand, Tibor Karaszi, Benjamin Nevarez, Simon Sabin, Darren Green, Allan Mitchell, Tony Rogerson, and many others—and to the Data Platform MVP lead, Jennifer Moser. This is a great program that I'm grateful for and proud to be part of. The level of expertise of this group is amazing, and I'm always excited when we all get to meet, both to share ideas and just to catch up at a personal level over beers.

Finally, to my students: Teaching about T-SQL is what drives me. It's my passion. Thanks for allowing me to fulfill my calling and for all the great questions that make me seek more knowledge.

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

<http://aka.ms/T-SQLFund3e/errata>

If you discover an error that is not already listed, please submit it to us at the same page.

If you need additional support, email Microsoft Press Book Support at mspinput@microsoft.com.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to <http://support.microsoft.com>.

Free ebooks from Microsoft Press

From technical overviews to in-depth information on special topics, the free ebooks from Microsoft Press cover a wide range of topics. These ebooks are available in PDF, EPUB, and Mobi for Kindle formats, ready for you to download at:

<http://aka.ms/mspressfree>

Check back often to see what is new!

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://aka.ms/tellpress>

We know you're busy, so we've kept it short with just a few questions. Your answers go directly to the editors at Microsoft Press. (No personal information will be requested.) Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>.

Background to T-SQL querying and programming

You're about to embark on a journey to a land that is like no other—a land that has its own set of laws. If reading this book is your first step in learning Transact-SQL (T-SQL), you should feel like Alice—just before she started her adventures in Wonderland. For me, the journey has not ended; instead, it's an ongoing path filled with new discoveries. I envy you; some of the most exciting discoveries are still ahead of you!

I've been involved with T-SQL for many years: teaching, speaking, writing, and consulting about it. For me, T-SQL is more than just a language—it's a way of thinking. In my first few books about T-SQL, I've written extensively on advanced topics, and for years, I have postponed writing about fundamentals. This is not because T-SQL fundamentals are simple or easy—in fact, it's just the opposite. The apparent simplicity of the language is misleading. I could explain the language syntax elements in a superficial manner and have you writing queries within minutes. But that approach would only hold you back in the long run and make it harder for you to understand the essence of the language.

Acting as your guide while you take your first steps in this realm is a big responsibility. I wanted to make sure that I spent enough time and effort exploring and understanding the language before writing about fundamentals. T-SQL is deep; learning the fundamentals the right way involves much more than just understanding the syntax elements and coding a query that returns the right output. You pretty much need to forget what you know about other programming languages and start thinking in terms of T-SQL.

Theoretical background

SQL stands for *Structured Query Language*. SQL is a standard language that was designed to query and manage data in relational database management systems (RDBMSs). An RDBMS is a database management system based on the relational model (a semantic model for representing data), which in turn is based on two mathematical branches: set theory and predicate logic. Many other programming languages and various aspects of computing evolved pretty much as a result of intuition. In contrast, to the degree that SQL is based on the relational model, it is based on a firm foundation—applied mathematics. T-SQL thus sits on wide and solid shoulders. Microsoft provides T-SQL as a dialect of, or extension to, SQL in Microsoft SQL Server data-management software, its RDBMS.

This section provides a brief theoretical background about SQL, set theory and predicate logic, the relational model, and types of database systems. Because this book is neither a mathematics book nor a design/data-modeling book, the theoretical information provided here is informal and by no means complete. The goals are to give you a context for the T-SQL language and to deliver the key points that are integral to correctly understanding T-SQL later in the book.

Language independence

The relational model is language independent. That is, you can apply data management and manipulation following the relational model's principles with languages other than SQL—for example, with C# in an object model. Today it is common to see RDBMSs that support languages other than a dialect of SQL, such as the CLR integration in SQL Server, with which you can handle tasks that historically you handled mainly with SQL, such as data manipulation.

Also, you should realize from the start that SQL deviates from the relational model in several ways. Some even say that a new language—one that more closely follows the relational model—should replace SQL. But to date, SQL is the industrial language used by all leading RDBMSs in practice.

See Also For details about the deviations of SQL from the relational model, as well as how to use SQL in a relational way, see this book on the topic: *SQL and Relational Theory: How to Write Accurate SQL Code*, Third Edition by C. J. Date (O'Reilly Media, 2015).

SQL

SQL is both an ANSI and ISO standard language based on the relational model, designed for querying and managing data in an RDBMS.

In the early 1970s, IBM developed a language called SEQUEL (short for Structured English QUery Language) for its RDBMS product called System R. The name of the language was later changed from SEQUEL to SQL because of a trademark dispute. SQL first became an ANSI standard in 1986, and then an ISO standard in 1987. Since 1986, the American National Standards Institute (ANSI) and the International Organization for Standardization (ISO) have been releasing revisions for the SQL standard every few years. So far, the following standards have been released: SQL-86 (1986), SQL-89 (1989), SQL-92 (1992), SQL:1999 (1999), SQL:2003 (2003), SQL:2006 (2006), SQL:2008 (2008), and SQL:2011 (2011). The SQL standard is made of multiple parts. Part 1 (Framework) and Part 2 (Foundation) pertain to the SQL language, whereas the other parts define standard extensions, such as SQL for XML and SQL-Java integration.

Interestingly, SQL resembles English and is also very logical. Unlike many programming languages, which use an imperative programming paradigm, SQL uses a declarative one. That is, SQL requires you to specify *what* you want to get and not *how* to get it, letting the RDBMS figure out the physical mechanics required to process your request.

SQL has several categories of statements, including Data Definition Language (DDL), Data Manipulation Language (DML), and Data Control Language (DCL). DDL deals with object definitions

and includes statements such as *CREATE*, *ALTER*, and *DROP*. DML allows you to query and modify data and includes statements such as *SELECT*, *INSERT*, *UPDATE*, *DELETE*, *TRUNCATE*, and *MERGE*. It's a common misunderstanding that DML includes only data-modification statements, but as I mentioned, it also includes *SELECT*. Another common misunderstanding is that *TRUNCATE* is a DDL statement, but in fact it is a DML statement. DCL deals with permissions and includes statements such as *GRANT* and *REVOKE*. This book focuses on DML.

T-SQL is based on standard SQL, but it also provides some nonstandard/proprietary extensions. Moreover, T-SQL does not implement all of standard SQL. In other words, T-SQL is both a subset and a superset of SQL. When describing a language element for the first time, I'll typically mention whether it is standard.

Set theory

Set theory, which originated with the mathematician Georg Cantor, is one of the mathematical branches on which the relational model is based. Cantor's definition of a set follows:

By a "set" we mean any collection M into a whole of definite, distinct objects m (which are called the "elements" of M) of our perception or of our thought.

—Joseph W. Dauben and Georg Cantor (Princeton University Press, 1990)

Every word in the definition has a deep and crucial meaning. The definitions of a set and set membership are axioms that are not supported by proofs. Each element belongs to a universe, and either is or is not a member of the set.

Let's start with the word *whole* in Cantor's definition. A set should be considered a single entity. Your focus should be on the collection of objects as opposed to the individual objects that make up the collection. Later on, when you write T-SQL queries against tables in a database (such as a table of employees), you should think of the set of employees as a whole rather than the individual employees. This might sound trivial and simple enough, but apparently many programmers have difficulty adopting this way of thinking.

The word *distinct* means that every element of a set must be unique. Jumping ahead to tables in a database, you can enforce the uniqueness of rows in a table by defining key constraints. Without a key, you won't be able to uniquely identify rows, and therefore the table won't qualify as a set. Rather, the table would be a *multiset* or a *bag*.

The phrase *of our perception or of our thought* implies that the definition of a set is subjective. Consider a classroom: one person might perceive a set of people, whereas another might perceive a set of students and a set of teachers. Therefore, you have a substantial amount of freedom in defining sets. When you design a data model for your database, the design process should carefully consider the subjective needs of the application to determine adequate definitions for the entities involved.

As for the word *object*, the definition of a set is not restricted to physical objects, such as cars or employees, but rather is relevant to abstract objects as well, such as prime numbers or lines.

What Cantor's definition of a set leaves out is probably as important as what it includes. Notice that the definition doesn't mention any order among the set elements. The order in which set elements are listed is not important. The formal notation for listing set elements uses curly brackets: {a, b, c}. Because order has no relevance, you can express the same set as {b, a, c} or {b, c, a}. Jumping ahead to the set of attributes (called *columns* in SQL) that make up the heading of a relation (called a *table* in SQL), an element is supposed to be identified by name—not by ordinal position.

Similarly, consider the set of tuples (called *rows* by SQL) that make up the body of the relation; an element is identified by its key values—not by position. Many programmers have a hard time adapting to the idea that, with respect to querying tables, there is no order among the rows. In other words, a query against a table can return table rows in *any order* unless you explicitly request that the data be sorted in a specific way, perhaps for presentation purposes.

Predicate logic

Predicate logic, whose roots reach back to ancient Greece, is another branch of mathematics on which the relational model is based. Dr. Edgar F. Codd, in creating the relational model, had the insight to connect predicate logic to both the management and querying of data. Loosely speaking, a *predicate* is a property or an expression that either holds or doesn't hold—in other words, is either true or false. The relational model relies on predicates to maintain the logical integrity of the data and define its structure. One example of a predicate used to enforce integrity is a constraint defined in a table called *Employees* that allows only employees with a salary greater than zero to be stored in the table. The predicate is "salary greater than zero" (T-SQL expression: *salary > 0*).

You can also use predicates when filtering data to define subsets, and more. For example, if you need to query the *Employees* table and return only rows for employees from the sales department, you use the predicate "department equals sales" in your query filter (T-SQL expression: *department = 'sales'*).

In set theory, you can use predicates to define sets. This is helpful because you can't always define a set by listing all its elements (for example, infinite sets), and sometimes for brevity it's more convenient to define a set based on a property. As an example of an infinite set defined with a predicate, the set of all prime numbers can be defined with the following predicate: "x is a positive integer greater than 1 that is divisible only by 1 and itself." For any specified value, the predicate is either true or not true. The set of all prime numbers is the set of all elements for which the predicate is true. As an example of a finite set defined with a predicate, the set {0, 1, 2, 3, 4, 5, 6, 7, 8, 9} can be defined as the set of all elements for which the following predicate holds true: "x is an integer greater than or equal to 0 and smaller than or equal to 9."

The relational model

The relational model is a semantic model for data management and manipulation and is based on set theory and predicate logic. As mentioned earlier, it was created by Dr. Edgar F. Codd, and later explained and developed by Chris Date, Hugh Darwen, and others. The first version of the relational model was proposed by Codd in 1969 in an IBM research report called "Derivability, Redundancy, and

Consistency of Relations Stored in Large Data Banks.” A revised version was proposed by Codd in 1970 in a paper called “A Relational Model of Data for Large Shared Data Banks,” published in the journal *Communications of the ACM*.

The goal of the relational model is to enable consistent representation of data with minimal or no redundancy and without sacrificing completeness, and to define data integrity (enforcement of data consistency) as part of the model. An RDBMS is supposed to implement the relational model and provide the means to store, manage, enforce the integrity of, and query data. The fact that the relational model is based on a strong mathematical foundation means that given a certain data-model instance (from which a physical database will later be generated), you can tell with certainty when a design is flawed, rather than relying solely on intuition.

The relational model involves concepts such as propositions, predicates, relations, tuples, attributes, and more. For nonmathematicians, these concepts can be quite intimidating. The sections that follow cover some key aspects of the model in an informal, nonmathematical manner and explain how they relate to databases.

Propositions, predicates, and relations

The common belief that the term *relational* stems from relationships between tables is incorrect. “Relational” actually pertains to the mathematical term *relation*. In set theory, a relation is a representation of a set. In the relational model, a relation is a set of related information, with the counterpart in SQL being a table—albeit not an exact counterpart. A key point in the relational model is that a single relation should represent a single set (for example, *Customers*). Note that operations on relations (based on relational algebra) result in a relation (for example, a join between two relations).



Note The relational model distinguishes between a *relation* and a *relation variable*, but to keep things simple, I won’t get into this distinction. Instead, I’ll use the term *relation* for both cases. Also, a relation is made of a heading and a body. The heading consists of a set of attributes (called *columns* in SQL), where each element is identified by an attribute name and a type name. The body consists of a set of tuples (called *rows* in SQL), where each element is identified by a key. To keep things simple, I’ll refer to a table as a set of rows.

When you design a data model for a database, you represent all data with relations (tables). You start by identifying propositions that you will need to represent in your database. A proposition is an assertion or a statement that must be true or false. For example, the statement, “Employee Itzik Ben-Gan was born on February 12, 1971, and works in the IT department” is a proposition. If this proposition is true, it will manifest itself as a row in a table of *Employees*. A false proposition simply won’t manifest itself. This presumption is known as the *closed-world assumption* (CWA).

The next step is to formalize the propositions. You do this by taking out the actual data (the body of the relation) and defining the structure (the heading of the relation)—for example, by creating predicates out of propositions. You can think of predicates as parameterized propositions. The heading of a relation comprises a set of attributes. Note the use of the term “set”; in the relational model,

attributes are unordered and distinct. An attribute is identified by an attribute name and a type name. For example, the heading of an *Employees* relation might consist of the following attributes (expressed as pairs of attribute names and type names): *employeeid* integer, *firstname* character string, *lastname* character string, *birthdate* date, and *departmentid* integer.

A *type* is one of the most fundamental building blocks for relations. A type constrains an attribute to a certain set of possible or valid values. For example, the type *INT* is the set of all integers in the range $-2,147,483,648$ to $2,147,483,647$. A type is one of the simplest forms of a predicate in a database because it restricts the attribute values that are allowed. For example, the database would not accept a proposition where an employee birth date is February 31, 1971 (not to mention a birth date stated as something like “abc!”). Note that types are not restricted to base types such as integers or character strings; a type also can be an enumeration of possible values, such as an enumeration of possible job positions. A type can be simple or complex. Probably the best way to think of a type is as a class—encapsulated data and the behavior supporting it. An example of a complex type is a geometry type that supports polygons.

Missing values

One aspect of the relational model is the source of many passionate debates—whether predicates should be restricted to two-valued logic. That is, in two-valued predicate logic, a predicate is either true or false. If a predicate is not true, it must be false. Use of two-valued predicate logic follows a mathematical law called “the law of excluded middle.” However, some say that there’s room for three-valued (or even four-valued) predicate logic, taking into account cases where values are missing. A predicate involving a missing value yields neither *true* nor *false*—it yields *unknown*.

Take, for example, a mobile phone attribute of an *Employees* relation. Suppose that a certain employee’s mobile phone number is missing. How do you represent this fact in the database? In a three-valued logic implementation, the mobile phone attribute should allow the use of a special marker for a missing value. Then a predicate comparing the mobile phone attribute with some specific number will yield *unknown* for the case with the missing value. Three-valued predicate logic refers to the three possible logical values that can result from a predicate—*true*, *false*, and *unknown*.

Some people believe that three-valued predicate logic is nonrelational, whereas others believe that it is relational. Codd actually advocated for four-valued predicate logic, saying that there were two different cases of missing values: missing but applicable (A-Values marker), and missing but inapplicable (I-Values marker). An example of “missing but applicable” is when an employee has a mobile phone, but you don’t know what the mobile phone number is. An example of “missing but inapplicable” is when an employee doesn’t have a mobile phone at all. According to Codd, two special markers should be used to support these two cases of missing values. SQL implements three-valued predicate logic by supporting the *NULL* marker to signify the generic concept of a missing value. Support for *NULLs* and three-valued predicate logic in SQL is the source of a great deal of confusion and complexity, though one can argue that missing values are part of reality. In addition, the alternative—using only two-valued predicate logic—is no less problematic.



Note As mentioned, a *NULL* is not a value but rather a marker for a missing value. Therefore, though unfortunately it's common, the use of the terminology "*NULL* value" is incorrect. The correct terminology is "*NULL* marker" or just "*NULL*." In the book, I chose to use the latter because it's more common in the SQL community.

Constraints

One of the greatest benefits of the relational model is the ability to define data integrity as part of the model. Data integrity is achieved through rules called *constraints* that are defined in the data model and enforced by the RDBMS. The simplest methods of enforcing integrity are assigning an attribute type with its attendant "nullability" (whether it supports or doesn't support *NULL*s). Constraints are also enforced through the model itself; for example, the relation *Orders*(*orderid*, *orderdate*, *duedate*, *shipdate*) allows three distinct dates per order, whereas the relations *Employees*(*empid*) and *EmployeeChildren*(*empid*, *childname*) allow zero to countable infinity children per employee.

Other examples of constraints include *candidate keys*, which provide entity integrity, and *foreign keys*, which provide referential integrity. A candidate key is a key defined on one or more attributes that prevents more than one occurrence of the same tuple (row in SQL) in a relation. A predicate based on a candidate key can uniquely identify a row (such as an employee). You can define multiple candidate keys in a relation. For example, in an *Employees* relation, you can define candidate keys on *employeeid*, on *SSN* (Social Security number), and others. Typically, you arbitrarily choose one of the candidate keys as the *primary key* (for example, *employeeid* in the *Employees* relation) and use that as the preferred way to identify a row. All other candidate keys are known as *alternate keys*.

Foreign keys are used to enforce referential integrity. A foreign key is defined on one or more attributes of a relation (known as the *referencing relation*) and references a candidate key in another (or possibly the same) relation. This constraint restricts the values in the referencing relation's foreign-key attributes to the values that appear in the referenced relation's candidate-key attributes. For example, suppose that the *Employees* relation has a foreign key defined on the attribute *departmentid*, which references the primary-key attribute *departmentid* in the *Departments* relation. This means that the values in *Employees.departmentid* are restricted to the values that appear in *Departments.departmentid*.

Normalization

The relational model also defines *normalization rules* (also known as *normal forms*). Normalization is a formal mathematical process to guarantee that each entity will be represented by a single relation. In a normalized database, you avoid anomalies during data modification and keep redundancy to a minimum without sacrificing completeness. If you follow Entity Relationship Modeling (ERM), and represent each entity and its attributes, you probably won't need normalization; instead, you will apply normalization only to reinforce and ensure that the model is correct. You can find the definition of ERM in the following Wikipedia article: https://en.wikipedia.org/wiki/Entity%E2%80%93relationship_model.

The following sections briefly cover the first three normal forms (1NF, 2NF, and 3NF) introduced by Codd.

1NF The first normal form says that the tuples (rows) in the relation (table) must be unique and attributes should be atomic. This is a redundant definition of a relation; in other words, if a table truly represents a relation, it is already in first normal form.

You achieve unique rows in SQL by defining a unique key for the table.

You can operate on attributes only with operations that are defined as part of the attribute's type. Atomicity of attributes is subjective in the same way that the definition of a set is subjective. As an example, should an employee name in an *Employees* relation be expressed with one attribute (*fullname*), two (*firstname* and *lastname*), or three (*firstname*, *middlename*, and *lastname*)? The answer depends on the application. If the application needs to manipulate the parts of the employee's name separately (such as for search purposes), it makes sense to break them apart; otherwise, it doesn't.

In the same way that an attribute might not be atomic enough based on the needs of the applications that use it, an attribute might also be subatomic. For example, if an address attribute is considered atomic for the applications that use it, not including the city as part of the address would violate the first normal form.

This normal form is often misunderstood. Some people think that an attempt to mimic arrays violates the first normal form. An example would be defining a *YearlySales* relation with the following attributes: *salesperson*, *qty2014*, *qty2015*, and *qty2016*. However, in this example, you don't really violate the first normal form; you simply impose a constraint—restricting the data to three specific years: 2014, 2015, and 2016.

2NF The second normal form involves two rules. One rule is that the data must meet the first normal form. The other rule addresses the relationship between nonkey and candidate-key attributes. For every candidate key, every nonkey attribute has to be fully functionally dependent on the entire candidate key. In other words, a nonkey attribute cannot be fully functionally dependent on part of a candidate key. To put it more informally, if you need to obtain any nonkey attribute value, you need to provide the values of all attributes of a candidate key from the same tuple. You can find any value of any attribute of any tuple if you know all the attribute values of a candidate key.

As an example of violating the second normal form, suppose that you define a relation called *Orders* that represents information about orders and order lines. (See Figure 1-1.) The *Orders* relation contains the following attributes: *orderid*, *productid*, *orderdate*, *qty*, *customerid*, and *companyname*. The primary key is defined on *orderid* and *productid*.

Orders	
PK	<u>orderid</u>
PK	<u>productid</u>
	orderdate qty customerid companyname

FIGURE 1-1 Data model before applying 2NF.

The second normal form is violated in Figure 1-1 because there are nonkey attributes that depend on only part of a candidate key (the primary key, in this example). For example, you can find the *orderdate* of an order, as well as *customerid* and *companyname*, based on the *orderid* alone.

To conform to the second normal form, you would need to split your original relation into two relations: *Orders* and *OrderDetails* (as shown in Figure 1-2). The *Orders* relation would include the attributes *orderid*, *orderdate*, *customerid*, and *companyname*, with the primary key defined on *orderid*. The *OrderDetails* relation would include the attributes *orderid*, *productid*, and *qty*, with the primary key defined on *orderid* and *productid*.

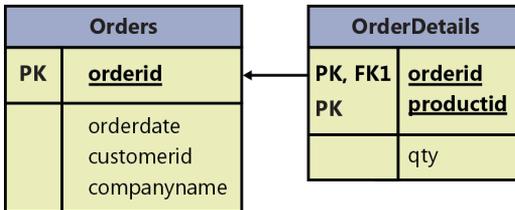


FIGURE 1-2 Data model after applying 2NF and before 3NF.

3NF The third normal form also has two rules. The data must meet the second normal form. Also, all nonkey attributes must be dependent on candidate keys nontransitively. Informally, this rule means that all nonkey attributes must be mutually independent. In other words, one nonkey attribute cannot be dependent on another nonkey attribute.

The *Orders* and *OrderDetails* relations described previously now conform to the second normal form. Remember that the *Orders* relation at this point contains the attributes *orderid*, *orderdate*, *customerid*, and *companyname*, with the primary key defined on *orderid*. Both *customerid* and *companyname* depend on the whole primary key—*orderid*. For example, you need the entire primary key to find the *customerid* representing the customer who placed the order. Similarly, you need the whole primary key to find the company name of the customer who placed the order. However, *customerid* and *companyname* are also dependent on each other. To meet the third normal form, you need to add a *Customers* relation (shown in Figure 1-3) with the attributes *customerid* (as the primary key) and *companyname*. Then you can remove the *companyname* attribute from the *Orders* relation.

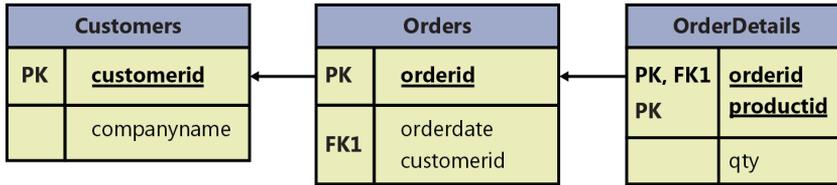


FIGURE 1-3 Data model after applying 3NF.

Informally, 2NF and 3NF are commonly summarized with the sentence, “Every non-key attribute is dependent on the key, the whole key, and nothing but the key—so help me Codd.”

There are higher normal forms beyond Codd’s original first three normal forms that involve compound primary keys and temporal databases, but they are outside the scope of this book.



Note SQL, as well as T-SQL, permit violating all the normal forms in real tables. It’s the data modeler’s prerogative and responsibility to design a normalized model.

Types of database systems

Two main types of systems, or workloads, use SQL Server as their database and T-SQL to manage and manipulate the data: online transactional processing (OLTP) and data warehouses (DWs). Figure 1-4 illustrates those systems and the transformation process that usually takes place between them.

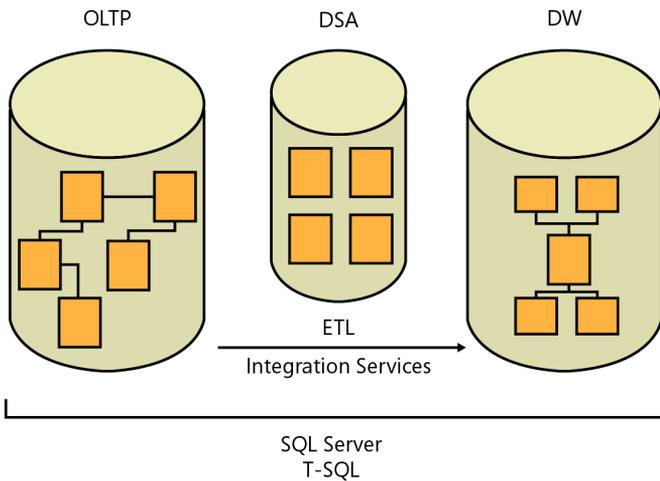


FIGURE 1-4 Classes of database systems.

Here's a quick description of what each acronym represents:

- OLTP: online transactional processing
- DSA: data-staging area
- DW: data warehouse
- ETL: extract, transform, and load

Online transactional processing

Data is entered initially into an online transactional processing system. The primary focus of an OLTP system is data entry and not reporting—transactions mainly insert, update, and delete data. The relational model is targeted primarily at OLTP systems, where a normalized model provides both good performance for data entry and data consistency. In a normalized environment, each table represents a single entity and keeps redundancy to a minimum. When you need to modify a fact, you need to modify it in only one place. This results in optimized performance for data modifications and little chance for error.

However, an OLTP environment is not suitable for reporting purposes because a normalized model usually involves many tables (one for each entity) with complex relationships. Even simple reports require joining many tables, resulting in complex and poorly performing queries.

You can implement an OLTP database in SQL Server and both manage it and query it with T-SQL.

Data warehouses

A *data warehouse* (DW) is an environment designed for data-retrieval and reporting purposes. When it serves an entire organization, such an environment is called a *data warehouse*; when it serves only part of the organization (such as a specific department) or a subject matter area in the organization, it is called a *data mart*. The data model of a data warehouse is designed and optimized mainly to support data-retrieval needs. The model has intentional redundancy, fewer tables, and simpler relationships, ultimately resulting in simpler and more efficient queries than an OLTP environment.

The simplest data-warehouse design is called a *star schema*. The star schema includes several dimension tables and a fact table. Each dimension table represents a subject by which you want to analyze the data. For example, in a system that deals with orders and sales, you will probably want to analyze data by dimensions such as customers, products, employees, and time.

In a star schema, each dimension is implemented as a single table with redundant data. For example, a product dimension could be implemented as a single *ProductDim* table instead of three normalized tables: *Products*, *ProductSubCategories*, and *ProductCategories*. If you normalize a dimension table, which results in multiple tables representing that dimension, you get what's known as a *snowflake dimension*. A schema that contains snowflake dimensions is known as a *snowflake schema*. A star schema is considered a special case of a snowflake schema.

The fact table holds the facts and measures, such as quantity and value, for each relevant combination of dimension keys. For example, for each relevant combination of customer, product,

employee, and day, the fact table would have a row containing the quantity and value. Note that data in a data warehouse is typically preaggregated to a certain level of granularity (such as a day), unlike data in an OLTP environment, which is usually recorded at the transaction level.

Historically, early versions of SQL Server mainly targeted OLTP environments, but eventually SQL Server also started targeting data-warehouse systems and data-analysis needs. You can implement a data warehouse as a SQL Server database and manage and query it with T-SQL.

The process that pulls data from source systems (OLTP and others), manipulates it, and loads it into the data warehouse is called *extract, transform, and load*, or *ETL*. SQL Server provides a tool called Microsoft SQL Server Integration Services (SSIS) to handle ETL needs.

Often the ETL process will involve the use of a data-staging area (DSA) between the OLTP and the DW. The DSA usually resides in a relational database, such as a SQL Server database, and is used as the data-cleansing area. The DSA is not open to end users.

SQL Server architecture

This section will introduce you to the SQL Server architecture, the different RDBMS flavors that Microsoft offers, the entities involved—SQL Server instances, databases, schemas, and database objects—and the purpose of each entity.

The ABCs of Microsoft RDBMS flavors

Initially, Microsoft offered mainly one enterprise-level RDBMS—an on-premises flavor called Microsoft SQL Server. These days, Microsoft offers an overwhelming plethora of options as part of its data platform, which constantly keeps evolving. Within its data platform, Microsoft offers three main RDBMS flavors, which you can think of as the *ABC flavors*: A for Appliance, B for Box (on-premises), and C for Cloud.

Box

The box, or on-premises RDBMS flavor, that Microsoft offers is called Microsoft SQL Server, or just SQL Server. This is the traditional flavor, usually installed on the customer's premises. The customer is responsible for everything—getting the hardware, installing the software, patching, high availability and disaster recovery, security, and everything else.

The customer can install multiple instances of the product on the same server (more on this in the next section) and can write queries that interact with multiple databases. It is also possible to switch the connection between databases, unless one of them is a contained database (defined later).

The querying language used is T-SQL. You can run all the code samples and exercises in this book on an on-premises SQL Server implementation, if you want. See the Appendix for details about obtaining and installing an evaluation edition of SQL Server, as well as creating the sample database.

Appliance

The idea behind the appliance flavor is to provide the customer a complete turnkey solution with preconfigured hardware and software. Speed is achieved by things being co-located, with the storage being close to the processing. The appliance is hosted locally at the customer site. Microsoft partners with hardware vendors such as Dell and HP to provide the appliance offering. Experts from Microsoft and the hardware vendor handle the performance, security, and availability aspects for the customer.

There are several appliances available today, one of which is the Microsoft Analytics Platform System (APS), which focuses on data warehousing and big data technologies. This appliance includes a data-warehouse engine called Parallel Data Warehouse (PDW), which implements massively parallel processing (MPP) technology. It also includes HDInsight, which is Microsoft's Hadoop distribution (big data solution). APS also includes a querying technology called PolyBase, which allows using T-SQL queries across relational data from PDW and nonrelational data from HDInsight.

Cloud

Cloud computing provides computing resources on demand from a shared pool of resources. Microsoft's RDBMS technologies can be provided both as private-cloud and public-cloud services. A *private cloud* is cloud infrastructure that services a single organization and usually uses virtualization technology. It's typically hosted locally at the customer site, and maintained by the IT group in the organization. It's about self-service agility, allowing the users to deploy resources on demand. It provides standardization and usage metering. The database engine is usually a box engine (hence the same T-SQL is used to manage and manipulate the data).

As for the public cloud, the services are provided over the network and available to the public. Microsoft provides two forms of public RDBMS cloud services: infrastructure as a service (IaaS) and platform as a service (PaaS). With IaaS, you provision a virtual machine (VM) that resides in Microsoft's cloud infrastructure. As a starting point, you can choose between several preconfigured VMs that already have a certain version and edition of SQL Server (box engine) installed on them. The hardware is maintained by Microsoft, but you're responsible for maintaining and patching the software. It's essentially like maintaining your own SQL Server installation—one that happens to reside on Microsoft's hardware.

With PaaS, Microsoft provides the database cloud platform as a service. It's hosted in Microsoft's data centers. Hardware, software installation and maintenance, high availability and disaster recovery, and patching are all responsibilities of Microsoft. The customer is still responsible for index and query tuning, however.

Microsoft provides a number of PaaS database offerings. For OLTP systems, it offers the Azure SQL Database service. It's also referred to more shortly as just SQL Database. The customer can have multiple databases on the cloud server (a conceptual server, of course) but cannot switch between databases.

Interestingly, Microsoft uses the same code base for SQL Database and SQL Server. So most of the T-SQL language surface is exposed (eventually) in both environments in the same manner. Therefore, most of the T-SQL you'll learn about in this book is applicable to both environments. You can read about the differences that do exist here: <https://azure.microsoft.com/en-us/documentation/articles/sql-database-transact-sql->

information. You should also note that the update and deployment rate of new versions of SQL Database are faster than that of the on-premises SQL Server. Therefore, some T-SQL features might be exposed in SQL Database before they show up in the on-premises SQL Server version.

Microsoft also provides a PaaS offering for data-warehouse systems called Microsoft Azure SQL Data Warehouse (also called Azure SQL Data Warehouse or just SQL Data Warehouse). This service is basically PDW/APS in the cloud. Microsoft uses the same code base for both the appliance and the cloud service. You manage and manipulate data in APS and SQL Data Warehouse with T-SQL, although it's not the same T-SQL surface as in SQL Server and SQL Database, yet.

Microsoft also offers other cloud data services, such as Data Lake for big data–related services, Azure DocumentDB for NoSQL document database services, and others.

Confused? If it's any consolation, you're not alone. Like I said, Microsoft provides an overwhelming plethora of database-related technologies. Curiously, the one thread that is common to many of them is T-SQL.

SQL Server instances

In the box product, an instance of SQL Server, as illustrated in Figure 1-5, is an installation of a SQL Server database engine or service. You can install multiple instances of on-premises SQL Server on the same computer. Each instance is completely independent of the others in terms of security and the data that it manages, and in all other respects. At the logical level, two different instances residing on the same computer have no more in common than two instances residing on two separate computers. Of course, same-computer instances do share the server's physical resources, such as CPU, memory, and disk.

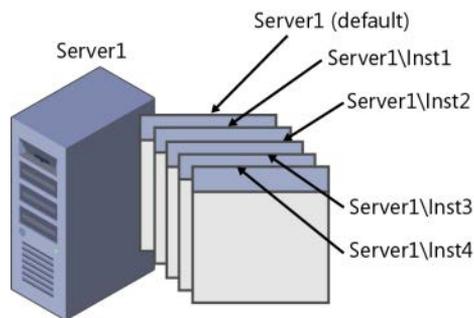


FIGURE 1-5 Multiple instances of SQL Server on the same computer.

You can set up one of the multiple instances on a computer as the *default instance*, whereas all others must be *named instances*. You determine whether an instance is the default or a named one upon installation; you cannot change that decision later. To connect to a default instance, a client application needs to specify the computer's name or IP address. To connect to a named instance, the client needs to specify the computer's name or IP address, followed by a backslash (\), followed by the instance name (as provided upon installation). For example, suppose you have two instances of SQL Server installed on a computer called *Server1*. One of these instances was installed as the default

instance, and the other was installed as a named instance called *Inst1*. To connect to the default instance, you need to specify only *Server1* as the server name. However, to connect to the named instance, you need to specify both the server and the instance name: *Server1\Inst1*.

There are various reasons why you might want to install multiple instances of SQL Server on the same computer, but I'll mention only a couple here. One reason is to save on support costs. For example, to test the functionality of features in response to support calls or reproduce errors that users encounter in the production environment, the support department needs local installations of SQL Server that mimic the user's production environment in terms of version, edition, and service pack of SQL Server. If an organization has multiple user environments, the support department needs multiple installations of SQL Server. Rather than having multiple computers, each hosting a different installation of SQL Server that must be supported separately, the support department can have one computer with multiple installed instances. Of course, you can achieve a similar result by using multiple virtual machines.

As another example, consider people like me who teach and lecture about SQL Server. For us, it is convenient to be able to install multiple instances of SQL Server on the same laptop. This way, we can perform demonstrations against different versions of the product, showing differences in behavior between versions, and so on.

As a final example, providers of database services sometimes need to guarantee their customers complete security separation of their data from other customers' data. At least in the past, the database provider could have a very powerful data center hosting multiple instances of SQL Server, rather than needing to maintain multiple less-powerful computers, each hosting a different instance. More recently, cloud solutions and advanced virtualization technologies make it possible to achieve similar goals.

Databases

You can think of a database as a container of objects such as tables, views, stored procedures, and other objects. Each instance of SQL Server can contain multiple databases, as illustrated in Figure 1-6. When you install an on-premises flavor of SQL Server, the setup program creates several system databases that hold system data and serve internal purposes. After the installation of SQL Server, you can create your own user databases that will hold application data.

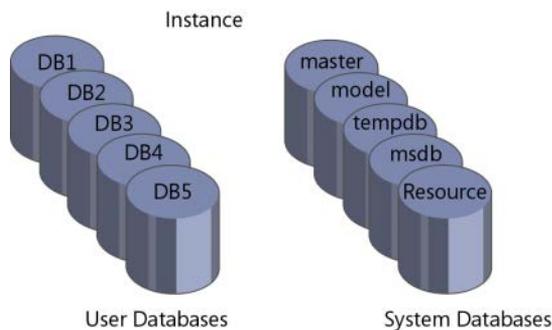


FIGURE 1-6 An example of multiple databases on a SQL Server instance.

The system databases that the setup program creates include *master*, *Resource*, *model*, *tempdb*, and *msdb*. A description of each follows:

- **master** The *master* database holds instance-wide metadata information, the server configuration, information about all databases in the instance, and initialization information.
- **Resource** The *Resource* database is a hidden, read-only database that holds the definitions of all system objects. When you query system objects in a database, they appear to reside in the *sys schema* of the local database, but in actuality their definitions reside in the *Resource* database.
- **model** The *model* database is used as a template for new databases. Every new database you create is initially created as a copy of *model*. So if you want certain objects (such as data types) to appear in all new databases you create, or certain database properties to be configured in a certain way in all new databases, you need to create those objects and configure those properties in the *model* database. Note that changes you apply to the *model* database will not affect existing databases—only new databases you create in the future.
- **tempdb** The *tempdb* database is where SQL Server stores temporary data such as work tables, sort and hash table data, row versioning information, and so on. With SQL Server, you can create temporary tables for your own use, and the physical location of those temporary tables is *tempdb*. Note that this database is destroyed and re-created as a copy of the *model* database every time you restart the instance of SQL Server.
- **msdb** The *msdb* database is used mainly by a service called SQL Server Agent to store its data. SQL Server Agent is in charge of automation, which includes entities such as jobs, schedules, and alerts. SQL Server Agent is also the service in charge of replication. The *msdb* database also holds information related to other SQL Server features, such as Database Mail, Service Broker, backups, and more.

In an on-premises installation of SQL Server, you can connect directly to the system databases *master*, *model*, *tempdb*, and *msdb*. In SQL Database, you can connect directly only to the system database *master*. If you create temporary tables or declare table variables (more on this topic in Chapter 11, “Programmable objects”), they are created in *tempdb*, but you cannot connect directly to *tempdb* and explicitly create user objects there.

You can create multiple user databases (up to 32,767) within an instance. A user database holds objects and data for an application.

You can define a property called *collation* at the database level that will determine default language support, case sensitivity, and sort order for character data in that database. If you do not specify a collation for the database when you create it, the new database will use the default collation of the instance (chosen upon installation).

To run T-SQL code against a database, a client application needs to connect to a SQL Server instance and be in the context of, or use, the relevant database. The application can still access objects from other databases by adding the database name as a prefix.

In terms of security, to be able to connect to a SQL Server instance, the database administrator (DBA) must create a *login* for you. The login can be tied to your Microsoft Windows credentials, in which case it is called a *Windows authenticated login*. With a Windows authenticated login, you can't provide login and password information when connecting to SQL Server because you already provided those when you logged on to Windows. The login can be independent of your Windows credentials, in which case it's called a *SQL Server authenticated login*. When connecting to SQL Server using a SQL Server authenticated login, you will need to provide both a login name and a password.

The DBA needs to map your login to a *database user* in each database you are supposed to have access to. The database user is the entity that is granted permissions to objects in the database.

SQL Server supports a feature called *contained databases* that breaks the connection between a database user and an instance-level login. The user (Windows or SQL authenticated) is fully contained within the specific database and is not tied to a login at the instance level. When connecting to SQL Server, the user needs to specify the database he or she is connecting to, and the user cannot subsequently switch to other user databases.

So far, I've mainly mentioned the logical aspects of databases. If you're using SQL Database, your only concern is that logical layer. You do not deal with the physical layout of the database data and log files, *tempdb*, and so on. But if you're using a box version of SQL Server, you are responsible for the physical layer as well. Figure 1-7 shows a diagram of the physical database layout.

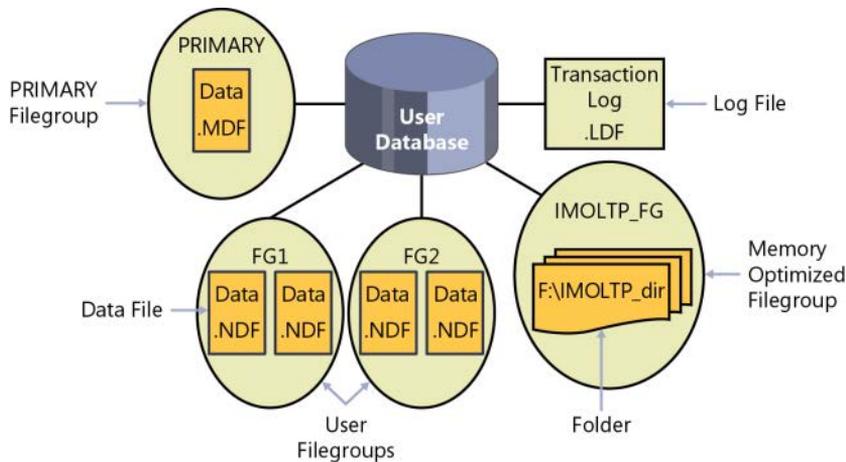


FIGURE 1-7 Database layout.

The database is made up of data files, transaction log files, and optionally checkpoint files holding memory-optimized data (part of a feature called *In-Memory OLTP*, which I describe shortly). When you create a database, you can define various properties for data and log files, including the file name, location, initial size, maximum size, and an autogrowth increment. Each database must have at least one data file and at least one log file (the default in SQL Server). The data files hold object data, and the log files hold information that SQL Server needs to maintain transactions.

Although SQL Server can write to multiple data files in parallel, it can write to only one log file at a time, in a sequential manner. Therefore, unlike with data files, having multiple log files does not result in a performance benefit. You might need to add log files if the disk drive where the log resides runs out of space.

Data files are organized in logical groups called *filegroups*. A filegroup is the target for creating an object, such as a table or an index. The object data will be spread across the files that belong to the target filegroup. Filegroups are your way of controlling the physical locations of your objects. A database must have at least one filegroup called *PRIMARY*, and it can optionally have other user filegroups as well. The *PRIMARY* filegroup contains the primary data file (which has an *.mdf* extension) for the database, and the database's system catalog. You can optionally add secondary data files (which have an *.ndf* extension) to *PRIMARY*. User filegroups contain only secondary data files. You can decide which filegroup is marked as the default filegroup. Objects are created in the default filegroup when the object creation statement does not explicitly specify a different target filegroup.

File extensions *.mdf*, *.ldf*, and *.ndf*

The database file extensions *.mdf* and *.ldf* are quite straightforward. The extension *.mdf* stands for *Master Data File* (not to be confused with the *master* database), and *.ldf* stands for *Log Data File*. According to one anecdote, when discussing the extension for the secondary data files, one of the developers suggested, humorously, using *.ndf* to represent "Not Master Data File," and the idea was accepted.

The SQL Server database engine includes a memory-optimized engine called In-Memory OLTP. You can use this feature to integrate memory-optimized objects, such as memory-optimized tables and natively compiled procedures, into your database. To do so, you need to create a filegroup in the database marked as containing memory-optimized data and, within it, at least one path to a folder. SQL Server stores checkpoint files with memory-optimized data in that folder, and it uses those to recover the data every time SQL Server is restarted.

Schemas and objects

When I said earlier that a database is a container of objects, I simplified things a bit. As illustrated in Figure 1-8, a database contains schemas, and schemas contain objects. You can think of a schema as a container of objects, such as tables, views, stored procedures, and others.

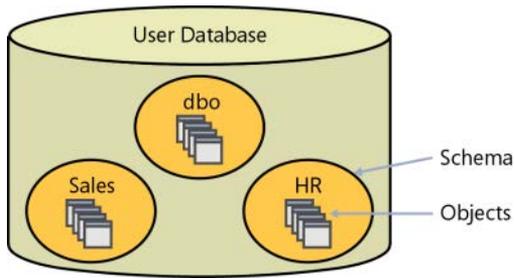


FIGURE 1-8 A database, schemas, and database objects.

You can control permissions at the schema level. For example, you can grant a user `SELECT` permissions on a schema, allowing the user to query data from all objects in that schema. So security is one of the considerations for determining how to arrange objects in schemas.

The schema is also a namespace—it is used as a prefix to the object name. For example, suppose you have a table named *Orders* in a schema named *Sales*. The schema-qualified object name (also known as the *two-part object name*) is *Sales.Orders*. You can refer to objects in other databases by adding the database name as a prefix (*three-part object name*), and to objects in other instances by adding the instance name as a prefix (*four-part object name*). If you omit the schema name when referring to an object, SQL Server will apply a process to resolve the schema name, such as checking whether the object exists in the user's default schema and, if the object doesn't exist, checking whether it exists in the *dbo* schema. Microsoft recommends that when you refer to objects in your code you always use the two-part object names. There are some relatively insignificant extra costs involved in resolving the schema name when you don't specify it explicitly. But as insignificant as this extra cost might be, why pay it? Also, if multiple objects with the same name exist in different schemas, you might end up getting a different object than the one you wanted.

Creating tables and defining data integrity

This section describes the fundamentals of creating tables and defining data integrity using T-SQL. Feel free to run the included code samples in your environment.



More Info If you don't know yet how to run code against SQL Server, the Appendix will help you get started.

As mentioned earlier, DML rather than DDL is the focus of this book. Still, you need to understand how to create tables and define data integrity. I won't go into the explicit details here, but I'll provide a brief description of the essentials.

Before you look at the code for creating a table, remember that tables reside within schemas, and schemas reside within databases. The examples use the book's sample database, *TSQV4*, and a schema called *dbo*.



More Info See the Appendix for details on creating the sample database.

The examples here use a schema named *dbo* that is created automatically in every database and is also used as the default schema for users who are not explicitly associated with a different schema.

Creating tables

The following code creates a table named *Employees* in the *dbo* schema in the *TSQLV4* database:

```
USE TSQLV4;

DROP TABLE IF EXISTS dbo.Employees;

CREATE TABLE dbo.Employees
(
    empid      INT          NOT NULL,
    firstname  VARCHAR(30) NOT NULL,
    lastname   VARCHAR(30) NOT NULL,
    hiredate   DATE         NOT NULL,
    mgrid      INT          NULL,
    ssn        VARCHAR(20) NOT NULL,
    salary     MONEY        NOT NULL
);
```

The *USE* statement sets the current database context to that of *TSQLV4*. It is important to incorporate the *USE* statement in scripts that create objects to ensure that SQL Server creates the objects in the specified database. In an on-premises SQL Server implementation, the *USE* statement can actually change the database context from one to another. In SQL Database, you cannot switch between different databases, but the *USE* statement will not fail as long as you are already connected to the target database. So even in SQL Database, I recommend having the *USE* statement to ensure that you are connected to the right database when creating your objects.

The *DROP IF EXISTS* command drops the table if it already exists. Note that this command was introduced in SQL Server 2016. If you're using earlier versions of SQL Server, use the following statement instead:

```
IF OBJECT_ID(N'dbo.Employees', N'U') IS NOT NULL DROP TABLE dbo.Employees;
```

The *IF* statement invokes the *OBJECT_ID* function to check whether the *Employees* table already exists in the current database. The *OBJECT_ID* function accepts an object name and type as inputs. The type *U* represents a user table. This function returns the internal object ID if an object with the specified input name and type exists, and *NULL* otherwise. If the function returns a *NULL*, you know that the object doesn't exist. In our case, the code drops the table if it already exists and then creates a new one. Of course, you can choose a different treatment, such as simply not creating the object if it already exists.

The *CREATE TABLE* statement is in charge of defining what I referred to earlier as the heading of the relation. Here you specify the name of the table and, in parentheses, the definition of its attributes (columns).

Notice the use of the two-part name *dbo.Employees* for the table name, as recommended earlier. If you omit the schema name, for ad-hoc queries SQL Server will assume the default schema associated with the database user running the code. For queries in stored procedures, SQL Server will assume the schema associated with the procedure's owner.

For each attribute, you specify the attribute name, data type, and whether the value can be *NULL* (which is called *nullability*).

In the *Employees* table, the attributes *empid* (employee ID) and *mgrid* (manager ID) are each defined with the *INT* (four-byte integer) data type; the *firstname*, *lastname*, and *ssn* (US Social Security number) are defined as *VARCHAR* (variable-length character string with the specified maximum supported number of characters); and *hiredate* is defined as *DATE* and *salary* is defined as *MONEY*.

If you don't explicitly specify whether a column allows or disallows *NULLs*, SQL Server will have to rely on defaults. Standard SQL dictates that when a column's nullability is not specified, the assumption should be *NULL* (allowing *NULLs*), but SQL Server has settings that can change that behavior. I recommend that you be explicit and not rely on defaults. Also, I recommend defining a column as *NOT NULL* unless you have a compelling reason to support *NULLs*. If a column is not supposed to allow *NULLs* and you don't enforce this with a *NOT NULL* constraint, you can rest assured that *NULLs* will occur. In the *Employees* table, all columns are defined as *NOT NULL* except for the *mgrid* column. A *NULL* in the *mgrid* column would represent the fact that the employee has no manager, as in the case of the CEO of the organization.

Coding style

You should be aware of a few general notes regarding coding style, the use of white spaces (space, tab, new line, and so on), and semicolons. I'm not aware of any formal coding styles. My advice is that you use a style that you and your fellow developers feel comfortable with. What ultimately matters most is the consistency, readability, and maintainability of your code. I have tried to reflect these aspects in my code throughout the book.

T-SQL lets you use white spaces quite freely in your code. You can take advantage of white space to facilitate readability. For example, I could have written the code in the previous section as a single line. However, the code wouldn't have been as readable as when it is broken into multiple lines that use indentation.

The practice of using a semicolon to terminate statements is standard and, in fact, is a requirement in several other database platforms. SQL Server requires the semicolon only in particular cases—but in cases where a semicolon is not required, using one doesn't cause problems. I recommend that you adopt the practice of terminating all statements with a semicolon. Not only will doing this improve the readability of your code, but in some cases it can save you some grief. (When a semicolon is required and is *not* specified, the error message SQL Server produces is not always clear.)



Note The SQL Server documentation indicates that not terminating T-SQL statements with a semicolon is a deprecated feature. This means that the long-term goal is to enforce use of the semicolon in a future version of the product. That's one more reason to get into the habit of terminating all your statements, even where it's currently not required.

Defining data integrity

As mentioned earlier, one of the great benefits of the relational model is that data integrity is an integral part of it. Data integrity enforced as part of the model—namely, as part of the table definitions—is considered *declarative data integrity*. Data integrity enforced with code—such as with stored procedures or triggers—is considered *procedural data integrity*.

Data type and nullability choices for attributes and even the data model itself are examples of declarative data integrity constraints. In this section, I will describe other examples of declarative constraints: primary key, unique, foreign key, check, and default constraints. You can define such constraints when creating a table as part of the *CREATE TABLE* statement, or you can define them for already-created tables by using an *ALTER TABLE* statement. All types of constraints except for default constraints can be defined as *composite constraints*—that is, based on more than one attribute.

Primary-key constraints

A primary-key constraint enforces the uniqueness of rows and also disallows *NULLs* in the constraint attributes. Each unique set of values in the constraint attributes can appear only once in the table—in other words, only in one row. An attempt to define a primary-key constraint on a column that allows *NULLs* will be rejected by the RDBMS. Each table can have only one primary key.

Here's an example of defining a primary-key constraint on the *empid* attribute in the *Employees* table that you created earlier:

```
ALTER TABLE dbo.Employees
  ADD CONSTRAINT PK_Employees
  PRIMARY KEY(empid);
```

With this primary key in place, you can be assured that all *empid* values will be unique and known. An attempt to insert or update a row such that the constraint would be violated will be rejected by the RDBMS and result in an error.

To enforce the uniqueness of the logical primary-key constraint, SQL Server will create a unique index behind the scenes. A *unique index* is a physical mechanism used by SQL Server to enforce uniqueness. Indexes (not necessarily unique ones) are also used to speed up queries by avoiding unnecessary full table scans (similar to indexes in books).

Unique constraints

A unique constraint enforces the uniqueness of rows, allowing you to implement the concept of alternate keys from the relational model in your database. Unlike with primary keys, you can define multiple unique constraints within the same table. Also, a unique constraint is not restricted to columns defined as *NOT NULL*.

The following code defines a unique constraint on the *ssn* column in the *Employees* table:

```
ALTER TABLE dbo.Employees
  ADD CONSTRAINT UNQ_Employees_ssn
  UNIQUE(ssn);
```

As with a primary-key constraint, SQL Server will create a unique index behind the scenes as the physical mechanism to enforce the logical unique constraint.

According to standard SQL, a column with a unique constraint is supposed to allow multiple *NULLs* (as if two *NULLs* were different from each other). However, SQL Server's implementation rejects duplicate *NULLs* (as if two *NULLs* were equal to each other). To emulate the standard unique constraint in SQL Server you can use a unique filtered index that filters only non-*NULL* values. For example, suppose that the column *ssn* allowed *NULLs*, and you wanted to create such an index instead of a unique constraint. You would have used the following code:

```
CREATE UNIQUE INDEX idx_ssn_notnull ON dbo.Employees(ssn) WHERE ssn IS NOT NULL;
```

The index is defined as a unique one, and the filter excludes *NULLs* from the index, so duplicate *NULLs* will be allowed, whereas duplicate non-*NULL* values won't be allowed.

Foreign-key constraints

A foreign-key enforces referential integrity. This constraint is defined on one or more attributes in what's called the *referencing* table and points to candidate-key (primary-key or unique-constraint) attributes in what's called the *referenced* table. Note that the referencing and referenced tables can be one and the same. The foreign key's purpose is to restrict the values allowed in the foreign-key columns to those that exist in the referenced columns.

The following code creates a table called *Orders* with a primary key defined on the *orderid* column:

```
DROP TABLE IF EXISTS dbo.Orders;

CREATE TABLE dbo.Orders
(
  orderid INT NOT NULL,
  empid INT NOT NULL,
  custid VARCHAR(10) NOT NULL,
  orderts DATETIME2 NOT NULL,
  qty INT NOT NULL,
  CONSTRAINT PK_Orders
  PRIMARY KEY(orderid)
);
```

Suppose you want to enforce an integrity rule that restricts the values supported by the *empid* column in the *Orders* table to the values that exist in the *empid* column in the *Employees* table. You can achieve this by defining a foreign-key constraint on the *empid* column in the *Orders* table pointing to the *empid* column in the *Employees* table, like the following:

```
ALTER TABLE dbo.Orders
  ADD CONSTRAINT FK_Orders_Employees
  FOREIGN KEY(empid)
  REFERENCES dbo.Employees(empid);
```

Similarly, if you want to restrict the values supported by the *mgrid* column in the *Employees* table to the values that exist in the *empid* column of the same table, you can do so by adding the following foreign key:

```
ALTER TABLE dbo.Employees
  ADD CONSTRAINT FK_Employees_Employees
  FOREIGN KEY(mgrid)
  REFERENCES dbo.Employees(empid);
```

Note that *NULLs* are allowed in the foreign-key columns (*mgrid* in the last example) even if there are no *NULLs* in the referenced candidate-key columns.

The preceding two examples are basic definitions of foreign keys that enforce a referential action called *no action*. No action means that attempts to delete rows from the referenced table or update the referenced candidate-key attributes will be rejected if related rows exist in the referencing table. For example, if you try to delete an employee row from the *Employees* table when there are related orders in the *Orders* table, the RDBMS will reject such an attempt and produce an error.

You can define the foreign key with actions that will compensate for such attempts (to delete rows from the referenced table or update the referenced candidate-key attributes when related rows exist in the referencing table). You can define the options *ON DELETE* and *ON UPDATE* with actions such as *CASCADE*, *SET DEFAULT*, and *SET NULL* as part of the foreign-key definition. *CASCADE* means that the operation (delete or update) will be cascaded to related rows. For example, *ON DELETE CASCADE* means that when you delete a row from the referenced table, the RDBMS will delete the related rows from the referencing table. *SET DEFAULT* and *SET NULL* mean that the compensating action will set the foreign-key attributes of the related rows to the column's default value or *NULL*, respectively. Note that regardless of which action you choose, the referencing table will have only orphaned rows in the case of the exception with *NULLs* that I mentioned earlier. Parents with no children are always allowed.

Check constraints

You can use a *check constraint* to define a predicate that a row must meet to be entered into the table or to be modified. For example, the following check constraint ensures that the salary column in the *Employees* table will support only positive values:

```
ALTER TABLE dbo.Employees
  ADD CONSTRAINT CHK_Employees_salary
  CHECK(salary > 0.00);
```

An attempt to insert or update a row with a nonpositive salary value will be rejected by the RDBMS. Note that a check constraint rejects an attempt to insert or update a row when the predicate evaluates to *FALSE*. The modification will be accepted when the predicate evaluates to either *TRUE* or *UNKNOWN*. For example, salary *-1000* will be rejected, whereas salaries *50000* and *NULL* will both be accepted (if the column allowed *NULLs*). As mentioned earlier, SQL is based on three-valued logic, which results in two actual actions. With a check constraint, the row is either accepted or rejected.

When adding check and foreign-key constraints, you can specify an option called *WITH NOCHECK* that tells the RDBMS you want it to bypass constraint checking for existing data. This is considered a bad practice because you cannot be sure your data is consistent. You can also disable or enable existing check and foreign-key constraints.

Default constraints

A default constraint is associated with a particular attribute. It's an expression that is used as the default value when an explicit value is not specified for the attribute when you insert a row. For example, the following code defines a default constraint for the *orderts* attribute (representing the order's time stamp):

```
ALTER TABLE dbo.Orders
  ADD CONSTRAINT DFT_Orders_orderts
  DEFAULT(SYSDATETIME()) FOR orderts;
```

The default expression invokes the *SYSDATETIME* function, which returns the current date and time value. After this default expression is defined, whenever you insert a row in the *Orders* table and do not explicitly specify a value in the *orderts* attribute, SQL Server will set the attribute value to *SYSDATETIME*.

When you're done, run the following code for cleanup:

```
DROP TABLE IF EXISTS dbo.Orders, dbo.Employees;
```

Conclusion

This chapter provided a brief background to T-SQL querying and programming. It presented a theoretical background, explaining the strong foundations that T-SQL is based on. It gave an overview of the SQL Server architecture and concluded with sections that demonstrated how to use T-SQL to create tables and define data integrity. I hope that by now you see that there's something special about SQL, and that it's not just a language that can be learned as an afterthought. This chapter equipped you with fundamental concepts—the actual journey is just about to begin.

This page intentionally left blank

Index

Symbols and Numbers

- (negative), 51–52
- (subtraction), 51–52
- % (percent) wildcard, 71–72
- (), parentheses, 51–52
- * (asterisk), SELECT clauses and, 41
- * (multiplication), 51–52
- ., = (comparison operators), 51–52
- .ldf (Log Data File), 18
- .mdf (Master Data File), 18
- .ndf (Not Master Data File), 18
- @@ identity, 255–257
- _ (underscore) wildcard, 72
- + (addition), 51–52
- + (concatenation), 51–52
- + (positive), 51–52
- = (assignment), 51–52
- [^character list or range] wildcard, 73
- [character-character] wildcard, 72–73
- [list of characters] wildcard, 72
- 1NF, 8
- 2NF, 8–10
- 3NF, 9–10

A

- ABC flavors (Appliance, Box, Cloud), 12–14
- accent sensitivity, character data, 61
- ACID (automaticity, consistency, isolation, and durability), 320–322
- addition (+), 51–52
- administrative tasks, dynamic SQL and, 379–383
- aggregate functions. *See also* COUNT; also pivoting data:
 - also window functions
 - aggregate window functions, 221–224
 - NULLs and, 33–35
 - running aggregates, subqueries and, 144–145

- aliases
 - common table expressions (CTEs), assigning in, 167–168
 - cross joins and, 104
 - derived tables and, 162–164
 - derived tables, nesting and, 165–166
 - inline table-valued functions (TVFs), 179
 - ORDER BY clause and, 42
 - SELECT clauses and, 36–41
 - self joins and, 106
- ALL, duplicates and, 194
- all-at-once operations, 58–59
- ALTER DATABASE, quoted identifiers, 62–63
- ALTER PROC, 385
- ALTER SEQUENCE, 259
- ALTER TABLE
 - defining data integrity, 22–25
 - sequence objects and, 261–262
- ALTER TABLE ADD CONSTRAINT, 264–266
- ALTER TABLE DROP CONSTRAINT, 264–266
- ALTER VIEW, 172
- American National Standards Institute (ANSI), 2
- AND
 - MERGE and, 275
 - precedence, 51–52
 - predicates and, 49
 - use of, 50
- ANSI (American National Standards Institute), 2
- Appliance, Box, Cloud (ABC flavors), 12–14
- application-time period tables, 297. *See also* temporal tables
- APPLY
 - table expressions, overview, 180–183
 - unpivoting data, 231–233
- arguments, derived table queries and, 164–165
- arithmetic operators, 50–52
- AS clause, 36–41
 - derived tables and, 161–162

assignment (=)

- assignment (=), 51–52
- asterisk (*), SELECT clauses and, 41
- AT TIME ZONE, 84–85
 - temporal table queries, 310
- attributes
 - filtering in outer joins, 118–119
 - normalization forms, 7–10
 - propositions, predicates, and relations, 5–6
- auditing, DDL triggers (data definition events), 388–390. See also temporal tables
- automaticity, transactions, 320–322
- automating administrative tasks, dynamic SQL and, 379–383
- autonumbering
 - assignment UPDATE and, 270–271
- AVG, 33–35. See also window functions
- Azure SQL Data Warehouse, 14
- Azure SQL Database
 - ABC flavors (Appliance, Box, Cloud), 13–14
 - global temporary tables, 376
 - isolation, overview, 332–333
 - locks and blocking, 322–323
 - READ COMMITTED SNAPSHOT, 343
 - row-versioning-based isolation, 339
 - source code, downloading of, 402–404
 - website link and access information, 395–396

B

- bag, 3
- batches
 - as unit of parsing, 364
 - as unit of resolution, 366
 - dynamic SQL, 379–383
 - overview of, 363–364
 - variables and, 361, 365
 - vs. transactions, 363
- BEGIN, 368
- BEGIN TRAN (TRANSACTION), 319–322
- BETWEEN
 - precedence, 51–52
 - temporal table queries, 308
 - use of, 49
- binding, batches and, 366
- bitemporal, 297. See also temporal tables
- blocking. See isolation, transactions; locks
- blocking chain, 328
- blocking_session_id, 330–331
- BULK INSERT, 253

C

- CACHE, sequence objects and, 258–259
- Cancel Executing Query, 331–332
- candidate keys. See also constraints
 - normalization forms, 8–9
 - overview of, 7
- Cantor, Georg, 3
- CASCADE, 24
- CASE expressions
 - overview of, 52–54
 - pivoting with grouped queries, 226–227
- case sensitivity, character data, 61
- CAST
 - correlated subqueries and, 141
 - date and time functions, 81–83
 - date and time literals, 77–78
- catalog views, metadata queries, 89–91
- CATCH, error handling, 390–394
- change management, DDL triggers (data definition events), 388–390
- CHAR, 50, 60–61
- character data
 - collation, 61–63
 - data types, overview, 60–61
 - quoted identifiers, 62–63
- CHARINDEX, 66
- CHECK
 - CASE expressions, overview of, 52–54
 - working with date and time separately, 78–79
- CHECK OPTION, views and, 177–178
- CHOOSE, CASE expressions, 54
- closed-world assumption, 5
- cloud computing, ABC flavors (Appliance, Box, Cloud), 13–14
- COALESCE, CASE expressions, 54
- Codd, Edgar F., 4–5
- code samples, resources for getting started
 - Azure SQL Database, getting started, 395–396
 - source code, downloading of, 402–404
 - SQL Server Books Online, 410–413
 - SQL Server installation, 396–401
 - SQL Server Management Studio, working with, 404–410
- coding style, T-SQL, 21–22
- COLLATE, 61–62
- collation, 16, 61–63
- column names
 - aliases, assigning in common table expressions (CTEs), 167–168

- derived tables and, 162
- identifier names, 30
- subquery substitution errors, 147–149
- COLUMNPROPERTY, 92
- COMMIT TRAN
 - lock modes and compatibility, 322–324
 - transactions, overview of, 319–322
- common table expressions (CTEs)
 - arguments, using, 167–168
 - column aliases, assigning, 167–168
 - defining multiple CTEs, 168–169
 - exercises, 183–188
 - exercises, solutions for, 188–192
 - multiple references in, 169
 - overview of, 167
 - recursive CTEs, 169–171
- comparison operators, 50–52
- composite joins, 110–111
 - exercises, 123–128
 - exercises, solutions for, 129–132
- COMPRESS, 70
- compression, character data and, 60–61
- CONCAT, 63–65
- concatenation (+), 51–52
- concurrency, transactions overview, 319–322
- conflict detection, SNAPSHOT and, 341–343
- consistency, transactions, 320–322. See also isolation, transactions
- CONSTRAINT, 264–266
- constraints
 - check constraints, defining, 24–25
 - default constraints, defining, 25
 - foreign-key constraints, defining, 23–24
 - overview of, 7
 - primary-key constraints, defining, 22
 - unique constraints, defining, 23
- contained databases, SQL Server architecture, 17
- CONTINUE, 369–370
- CONVERT
 - date and time functions, 81–83
 - date and time literals, 77–78
 - working with date and time separately, 79
- correlated subqueries. See also subqueries
 - EXISTS and, 141–143
 - overview of, 139–141
- COUNT. See also window functions
 - GROUP BY and, 33–35
 - outer joins, use with, 121–123
- CREATE DATABASE, DDL triggers (data definition events), 388–390

- CREATE DEFAULT, 365
- CREATE FUNCTION, 365
- CREATE PROCEDURE, 365
- CREATE RULE, 365
- CREATE SCHEMA, 365
- CREATE SEQUENCE, 258–259
- CREATE TABLE
 - DDL triggers (data definition events), 388–390
 - defining data integrity, 22–25
 - use of, 20
- CREATE TRIGGER, 365
- CREATE VIEW, 365
- CROSS APPLY
 - JOIN and, 180
 - overview of, 180–183
 - unpivoting data, 232–233
- cross-joins
 - CROSS APPLY and, 180
 - exercises, 123–128
 - exercises, solutions for, 129–132
 - overview of, 103–107
- CUBE, 235–236
- CURRENT_TIMESTAMP, 80–81
- cursor, 42
 - as programmable objects, 370–373
- custom sequence, assignment UPDATE and, 270–271
- CYCLE, sequence objects and, 258–259

D

- Darwen, Hugh, 4–5
- data definition events (DDL triggers), 387–390
- data definition language (DDL), batches and, 366
- data integrity
 - constraints and, 7
 - defining, 22–25
- Data Lake, 14
- data manipulation events (DML triggers), 387–388
- Data Manipulation Language (DML), 249
 - batches and, 366
 - nested DML, 285–287
- data marts, overview of, 11–12
- data types
 - character data, 60–61
 - N prefix, use of, 50
 - operands, precedence of, 50–52
- data warehouses (DWs), 10–12
- database system types, 10–12
- database user, SQL Server logins, 17

DATABASEPROPERTYEX

- DATABASEPROPERTYEX, 92
- databases, SQL Server
 - architecture overview, 15–18
 - file extensions (.mdf, .ldf, .ndf), 18
 - schemas and objects, 18–19
- DATALength, 65–66
- DATE
 - data types, 73–74
 - literals, 74–78
 - working with date and time separately, 78–79
- date and time data
 - aggregate window functions, 223–224
 - AT TIME ZONE, 84–85
 - CAST, CONVERT, and PARSE, 81–83
 - current date and time functions, 80–81
 - DATEADD, 85
 - DATEDIFF and DATEDIFF_BIG, 86–87
 - DATENAME, 88
 - DATEPART, 87
 - EOMONTH, 89
 - filtering date ranges, 79–80
 - FROMPARTS, 88–89
 - ISDATE, 88
 - literals, 74–78
 - missing values, in outer joins, 116–118
 - SWITCHOFFSET, 83
 - types of, 73–74
 - working with separately, 78–79
 - YEAR, MONTH, and DAY, 87–88
- Date, Chris, 4–5
- DATEADD, 85
- DATEDIFF and DATEDIFF_BIG, 86–87
- DATEFORMAT, 75
- DATENAME, 88
- DATEPART, 87
- DATETIME
 - data types, 73–74
 - literals, 74–78
 - working with date and time separately, 78–79
- DATETIME2
 - data types, 73–74
 - literals, 74–78
 - temporal tables, creating, 298–301
- DATETIMEOFFSET
 - data types, 73–74
 - literals, 74–78
 - SWITCHOFFSET and, 83
- DAY, 87–88
- DBCC CHECKIDENT, 257
- DDL triggers (data definition events), 387–390
- DEADLOCK_PRIORITY, 345
- deadlocks, 337, 345–348
- deadly embrace deadlock, 347
- declarative data integrity, 22–25
- DECLARE
 - temporary table variables, 377–378
 - variables and, 361
- DECOMPRESS, 70
- default constraints, defining, 25
- DELETE
 - DML triggers (data manipulation events), 387–388
 - exercises, 287–291
 - exercises, solutions for, 291–295
 - joins and, 265–266
 - OUTPUT and, 282–283
 - overview of, 262–266
 - row-versioning isolation and, 339
 - table expressions, modifying data with, 276–278
 - views and, 172
- DENSE_RANK, window function ranking, 216–219
- derived tables
 - arguments, use of, 164–165
 - column aliases, assigning, 163–164
 - exercises, 183–188
 - exercises, solutions for, 188–192
 - multiple references, 166
 - nesting, 165–166
 - overview of, 161–162
- DESC, 43
- dictionary sorting, character data, 61
- dirty reads. See isolation, transactions
- DISTINCT
 - aggregate functions and, 34–35
 - duplicates and, 39–41, 194
 - EXCEPT (DISTINCT), 199–200
 - INTERSECT (DISTINCT), 196–197
 - multivalued subqueries and, 138
 - ORDER BY clause and, 43
 - UNION (DISTINCT), 194–196
 - window functions, ranking of, 218–219
- distinct, defined, 3
- DML. See Data Manipulation Language (DML)
- DML triggers (data manipulation events), 387–388
- DROP CONSTRAINT, sequence objects and, 261–262
- DROP IF EXISTS, table creation, 20
- duplicates, DISTINCT and, 39–41, 194
- durability, transactions, 320–322
- DWs (data warehouses), 10–12
- dynamic SQL, 379–383

E

ELSE

CASE expressions, overview of, 52–54

IF . . . ELSE flow element, 367–368

ELSE NULL, CASE expressions, 52–54

ENCRYPTION, views and, 175–176

END, 368

Entity Relationship Modeling (ERM), 7

EOMONTH, 89

equi and non-equi joins, 111–112

exercises, 123–128

exercises, solutions for, 129–132

error handling

deadlocks and, 347–348

programmable objects and, 390–394

stored procedures and, 385

transactions, overview of, 320–322

ERROR_LINE, 391

ERROR_MESSAGE, 391

ERROR_NUMBER, 391

ERROR_PROCEDURE, 391

ERROR_SEVERITY, 391

ERROR_STATE, 391

ESCAPE character, 73

ETL (extract, transform, and load), 12

EVENTDATA, DDL triggers (data definition events),

388–390

EXCEPT

exercises, 204–208

exercises, solutions for, 208–211

precedence and, 201–202

use of, 199–201

EXCEPT (DISTINCT), 199–200

EXCEPT ALL, 199–201

exclusive locks, 322–324

EXEC

dynamic SQL, 379–381

INSERT EXEC, 251–252

execution plan, SQL server, 380–381

exercises

DELETE, 287–291

DELETE, solutions for, 291–295

grouping set, 241–244

grouping set, solutions for, 245–248

INSERT, 287–291

INSERT, solutions for, 291–295

JOIN, 123–128

JOIN, solutions for, 129–132

MERGE, 287–291

MERGE, solutions for, 291–295

OFFSET-FETCH, 287–291

OFFSET-FETCH, solutions for, 291–295

OUTPUT, 287–291

OUTPUT, solutions for, 291–295

pivoting and unpivoting data, 241–244

pivoting and unpivoting data, solutions for,
245–248

SELECT statements, 93–97

SELECT statements, solutions for, 97–102

subqueries, 150–154

subqueries, solutions for, 154–159

table expressions, 183–188, 287–291

table expressions, solutions for, 188–192, 291–295

temporal tables, 311–313

temporal tables, solutions for, 313–317

TOP, 287–291

TOP, solutions for, 291–295

transactions, isolation and locks, 348–359

UPDATE, 287–291

UPDATE, solutions for, 291–295

windows functions, 241–244

windows functions, solutions for, 245–248

exercises, resources for getting started

Azures SQL Database, getting started, 395–396

source code, downloading of, 402–404

SQL Server Books Online, 410–413

SQL Server installation, 396–401

SQL Server Management Studio, working with,
404–410

EXISTS

correlated subqueries and, 141–143

subqueries, NULL trouble and, 147

extract, transform, and load (ETL), 12

F

FALSE

check constraints and, 24–25

IF . . . ELSE flow element, 367–368

NULLs, overview of, 54–58

subqueries, NULL trouble and, 146–147

WHERE and, 32

WHILE flow element, 368–370

filters. See also subqueries

date ranges, 79–80

deadlocks, avoiding, 347

HAVING, overview of, 35–36

OFFSET-FETCH, overview of, 46–47

FIRST_VALUE

- filters. See also subqueries *continued*
 - predicate logic, overview of, 4
 - TOP filters, overview of, 44–46
 - WHERE, overview of, 31–32
 - WITH TIES and, 46
- FIRST_VALUE, offset window functions, 219–221
- flow elements, 367–370
 - IF . . . ELSE flow element, 367–368
 - WHILE flow element, 368–370
- fn_helpcollations, 61–63
- FOR SYSTEM_TIME CONTAINED IN, 308–310
- FOR SYSTEM_TIME, temporal table queries, 304–310
- FOR XML, views and ORDER BY clause, 173–174
- foreign-key constraints. See also constraints
 - data integrity, defining, 23–24
 - overview of, 7
- FORMAT, 69
- FROM. See also JOIN
 - derived tables and, 161–162
 - logical query processing order, 28
 - overview of, 29–30
- FROMPARTS, 88–89
- functions
 - CHARINDEX, 66
 - COMPRESS and DECOMPRESS, 70
 - CONCAT, string concatenation, 63–65
 - date and time functions, 80–89
 - FORMAT, 69
 - LEFT and RIGHT, 65
 - LEN and DATALENGTH, 65–66
 - PATINDEX, 66
 - REPLACE, 66–67
 - REPLICATE, 67–68
 - RTRIM and LTRIM, 69
 - STRING_SPLIT, 71
 - STUFF, 68
 - SUBSTRING, 65
 - UPPER and LOWER, 68–69

G

- generate globally unique identifiers (GUIDs), INSERT SELECT and, 251
- GENERATED ALWAYS AS ROW END, 298–301
- GENERATED ALWAYS AS ROW START, 298–301
- GETDATE, 80–81
- GETUTCDATE, 80–81
- global temporary tables, 376–377
- GO, 363–367

- GROUP BY. See also grouping sets
 - CUBE, 236
 - derived tables, assigning aliases, 163–164
 - GROUPING SETS, 235–236
 - HAVING and, 35–36
 - logical query processing order, 28
 - overview of, 32–35
 - pivoting with grouped queries, 226–227
 - ROLLUP, 236–237
- GROUPING, 235, 238–240
- grouping sets. See also GROUP BY
 - CUBE, 236
 - exercises, 241–244
 - exercises, solutions for, 245–248
 - GROUPING and GROUPING_ID functions, 238–240
 - GROUPING SETS subclause, 235–236
 - overview, 234–235
 - ROLLUP, 236–237
- GROUPING SETS, 235–236
- GROUPING_ID, 235, 238–240
- GZIP algorithm, 70

H

- HAVING
 - CASE expressions, overview of, 52–54
 - GROUP BY and, 33
 - logical query processing order, 28
 - overview of, 35–36
- HDInsight, 13
- historical data. See temporal tables
- HOLDLOCK, 332

I

- IaaS (infrastructure as a service)
 - ABC flavors (Appliance, Box, Cloud), 13–14
- IDENT_CURRENT, 255–257
- identifier names
 - delimiting of, 30
 - quoted identifiers, 62–63
- identity property, 254–257
- IDENTITY_INSERT, 256–257
- IF . . . ELSE flow element, 367–368
- IIF, CASE expressions, 54
- IMPLICIT_TRANSACTIONS, 319–320
- IN
 - multivalued subquery examples, 136–139
 - precedence, 51–52

- subqueries, NULL trouble and, 146–147
- use of, 49
- inconsistent analysis, 335, 344–345
- INCREMENT BY, sequence objects and, 258–259
- information schema views, metadata queries, 91
- INFORMATION_SCHEMA, 91
- infrastructure as a service (IaaS), 13–14
- inline table-valued functions (ITVFs)
 - exercises, 183–188
 - exercises, solutions for, 188–192
 - overview of, 171–172, 178–180
- In-Memory OLTP, 18
- inner joins, 107–109
 - exercises, 123–128
 - exercises, solutions for, 129–132
- INSERT
 - DML triggers (data manipulation events), 387–388
 - exercises, 287–291
 - exercises, solutions for, 291–295
 - identity property and, 254–257
 - MERGE and, 274
 - OUTPUT and, 280–282
 - row-versioning isolation and, 339
 - table expressions, modifying data with, 276–278
 - views and, 172
- INSERT EXEC, 251–252
- INSERT SELECT, 251
- INSERT VALUES, 249–251
- INT, 51
- integer sequences, cross joins and, 106–107
- International Organization for Standardization (ISO), 2
- INTERSECT
 - exercises, 204–208
 - exercises, solutions for, 208–211
 - precedence and, 201–202
 - use of, 196–199
- INTERSECT (DISTINCT), 196–197
- INTERSECT ALL, 196–199
- INTO, SELECT INTO, 252–253
- IS NOT NULL, 55–58
- IS NULL
 - CASE expressions and, 54
 - overview of, 55–58
- ISDATE, 88
- ISO (International Organization for Standardization), 2
- ISO/ANSI SQL-89 syntax
 - cross joins, 105
 - inner joins, 108–109
- ISO/ANSI SQL-92 syntax
 - cross joins, 104

- inner joins, 107–108
- isolation, transactions, 320–322
 - deadlocks, 345–348
 - exercises, 348–359
 - overview of, 332–333
 - READ COMMITTED, 334–337
 - READ COMMITTED SNAPSHOT, 343–345
 - READ UNCOMMITTED, 333–334
 - REPEATABLE READ, 335–337
 - SERIALIZABLE, 337–338
 - SNAPSHOT, 339–343
 - SNAPSHOT, conflict detection, 341–343

J

- JOIN
 - composite joins, 110–111
 - CROSS APPLY and, 180
 - cross-joins, 103–107
 - DELETE and, 265–266
 - derived tables, multiple references, 166
 - exercises, 123–128
 - exercises, solutions for, 129–132
 - inner joins, 107–109
 - multi-join queries, 113
 - natural joins, 111
 - non-equi joins, 111–112
 - OUTER APPLY and, 181
 - outer joins
 - COUNT aggregate, use of, 121–123
 - filtering attributes from nonpreserved side, 118–119
 - fundamentals of, 113–116
 - in multi-join queries, 119–121
 - missing values, inclusion of, 116–118
 - overview of, 103
 - self joins, 105–106
 - UPDATE and, 268–270
 - vs. subqueries, 137

K

- keys, SQL queries and, 39–41

L

- LAG, 378
 - offset window functions, 219–221
- language independence, 2

language standards organizations

- language standards organizations, 2
- LANGUAGE/DATEFORMAT, 75–76
- LAST_VALUE, offset window functions, 219–221
- LEAD, offset window functions, 219–221
- LEFT, 65
- LEN, 65–66
- LIKE
 - precedence, 51–52
 - use of, 50, 71–73
- linked history tables, 298–301
- literals, date and time data, 74–78
- local temporary tables, 374–375
- LOCK_ESCALATION, 325
- LOCK_TIMEOUT, 331–332
- locks. See also isolation, transactions
 - deadlocks, 337, 345–348
 - exercises, 348–359
 - lockable resource types, 324–325
 - modes and compatibility, 322–324
 - troubleshooting, 325–332
- logical operators, 50
- login, SQL Server databases, 17
- loops, WHILE element and, 368–370
- LOWER, 68–69
- LTRIM, 69

M

- massively parallel processing (MPP), 13
- master database, SQL Server architecture, 16–18
- MAX, 33–35, 61
- MAXVALUE, sequence objects and, 258–259
- MERGE
 - DML triggers (data manipulation events), 387–388
 - exercises, 287–291
 - exercises, solutions for, 291–295
 - merging data overview, 271–275
 - OUTPUT and, 284–285
 - table expressions, modifying data with, 276–278
- metadata, querying of, 89–92
- Microsoft Analytics Platform System (APS), 13–14
- Microsoft Azure SQL Data Warehouse, 14
- Microsoft Azure SQL Database. See Azure SQL Database
- Microsoft SQL Server. See SQL Server
- MIN, 33–35
- MINVALUE, sequence objects and, 258–259
- missing values, 6–7
 - in outer joins, 116–118

- model database, SQL Server architecture, 16–18
- MONTH, 79–80, 87–88
- msdb database, SQL Server architecture, 16–18
- MTD, aggregate window functions, 223–224
- multi-join queries
 - exercises, 123–128
 - exercises, solutions for, 129–132
 - outer joins in, 119–121
 - use of, 113
- multiplication (*), 51–52
- multiset, 3
- multi-statement table-valued function (TVF), 178
- multivalued subqueries
 - defined, 133
 - self-contained subquery examples, 136–139

N

- N (National), data type prefix, 50, 60–61
- names, schemas and objects, 18–19. See also aliases
- natural joins, 111
- NCHAR, 50, 60–61
- negative (-), 51–52
- nested DML, 285–287
- nested queries, 133
 - derived tables and, 165–166
- NEWID, 251, 383
- NEXT VALUE FOR, 259–262
- next values, subquery returns, 143–144
- no action, foreign-key constraints, 24
- NOLOCK, 332
- non-equi joins. See equi and non-equi joins
- nonkey attributes, normalization forms, 8–10
- nonrepeatable reads, 335, 344–345
- normal forms. See normalization
- normalization, overview, 7–10
- NOT
 - combining logical expressions, 50
 - multivalued subqueries and, 137
 - precedence, 51–52
 - subqueries, NULL trouble and, 146–147
- NOT EXISTS, NULL trouble and, 147
- NOT IN, NULL trouble and, 146–147
- NOT NULL, grouping and, 238–240
- NTILE, window function ranking, 216–219
- NULL
 - aggregate functions and, 34–35
 - CASE expressions and, 226
 - data integrity and, 22–25

DISTINCT and duplicates, 194
 GROUPING and GROUPING_ID functions and, 238–240
 IF . . . ELSE flow element, 367–368
 INSERT VALUES and, 250
 INTERSECT and, 197
 missing values, overview, 6–7
 multivalued subqueries and, 137
 outer joins, filtering attributes, 118–119
 outer joins, fundamentals of, 115–116
 outer joins, in multi-join queries, 119–121
 overview of, 54–58
 scalar subqueries and, 135
 SET NULL and, 24
 string concatenation and, 63–65
 subqueries, troubles with, 145–147
 tables, creating, 21
 variable values and, 363
 NUMERIC, 51
 NVARCHAR, 50, 60–61

O

object names, schema qualifying of, 29
 object, defined, 3
 OBJECT_ID, table creation, 20
 OBJECTPROPERTY, 92
 objects, SQL Server architecture, 18–19. See also programmable objects
 offset window functions, 219–221
 OFFSET-FETCH
 data modification with, 278–280
 derived tables and, 162
 exercises, 287–291
 exercises, solutions for, 291–295
 overview of, 46–47
 views and ORDER BY clause, 173–174
 OLTP (online transactional processing) databases, 10–12
 ON
 outer joins, fundamentals of, 115–116
 outer joins, in multi-join queries, 119–121
 ON DELETE CASCADE, 24
 online transactional processing (OLTP) databases, overview, 10–12
 operators. See also APPLY; also JOIN; also PIVOT; also UNPIVOT
 compound operators, UPDATE and, 267–268
 overview of, 49–52
 precedence rules, 51–52
 string concatenation, 63–65
 OR
 combining logical expressions, 50
 precedence, 51–52
 predicates and, 49
 ORDER BY
 CASE expressions, overview of, 52–54
 circumventing unsupported logical phrases, 203–204
 derived tables and, 162
 GROUP BY and, 33
 INTERSECT ALL and, 197–199
 logical query processing order, 28
 offset window functions and, 220–221
 OFFSET-FETCH, overview of, 46–47
 overview of, 41–43
 set operators and, 193
 TOP filters and, 45–46
 views, table expressions, 172–174
 window functions, overview of, 47–49, 214–215
 ordinal positions, ORDER BY clause and, 43
 OUTER APPLY
 JOIN and, 181
 overview of, 180–183
 outer joins
 COUNT aggregate, use of, 121–123
 exercises, 123–128
 exercises, solutions for, 129–132
 filtering attributes from nonpreserved side, 118–119
 fundamentals of, 113–116
 in multi-join queries, 119–121
 missing values, inclusion of, 116–118
 outer queries, 133. See also subqueries; also table expressions
 OUTPUT
 DELETE and, 282–283
 exercises, 287–291
 exercises, solutions for, 291–295
 INSERT and, 280–282
 MERGE and, 284–285
 nested DML, 285–287
 overview of, 280
 UPDATE and, 283–284
 OVER
 empty parens () and, 221–222
 sequence objects and, 260–262
 window functions, overview of, 47–49, 213–216

P

PaaS (platform as a service)

ABC flavors (Appliance, Box, Cloud), 13–14

Parallel Data Warehouse (PDW), 13

parentheses, (), 51–52

PARSE

date and time functions, 81–83

date and time literals, 77–78

parsing, batches and, 364

PARTITION BY

INTERSECT ALL and, 197–199

window functions, overview of, 47–49, 215

partitions, offset window functions, 219–221

PATINDEX, 66

percent (%) wildcard, 71–72

PERCENT, TOP filters and, 45

performance

dynamic SQL and, 379–383

object names, schema qualifying, 29

query filters and, 32

row-versioning isolation and, 339

sp_executesql stored procedure, 380–381

SQL Server architecture, physical layout, 17–18

stored procedures and, 385

PERIOD FOR SYSTEM_TIME, temporal table creation, 298–301

permissions, database schemas and objects, 18–19

phantom reads, 337–338

SNAPSHOT and, 339

PIVOT

dynamic SQL and, 379, 381–383

pivoting data, overview, 227–229

pivoting data

exercises, 241–244

exercises, solutions for, 245–248

overview of, 224–226

unpivoting data, 230–234

with PIVOT operator, 227–229

within grouped queries, 226–227

platform as a service (PaaS), 13–14

point-in-time analysis. See temporal tables

policy enforcement, DDL triggers (data definition events), 388–390

PolyBase, 13

positive (+), 51–52

predicates. See also specific predicate names

LIKE predicate, 71–73

NULLs, overview of, 54–58

overview of, 49–52

predicate logic, 4

relational model, overview, 5–6

previous table versions. See temporal tables

previous values, subquery returns, 143–144

primary-key constraints, data integrity, 22

private cloud, ABC flavors (Appliance, Box, Cloud), 13–14

procedural data integrity, 22–25

programmable objects

batches and, 363–367

cursors, 370–373

dynamic SQL, 379–383

error handling, 390–394

flow element, 367–370

routines, 383

stored procedures, 385–386

triggers, 387–390

user-defined functions, 383–384

temporary tables, 374–379

variables, 361–363

propositions, relational model overview, 5–6

public cloud, ABC flavors (Appliance, Box, Cloud), 13–14

Q

queries. See also SELECT statement; also subqueries;

also table expressions

aliases, use of, 36–41

all-at-once operations, 58–59

CASE expressions, overview of, 52–54

character data, working with, 60–63

date and time data, working with, 73–89

exercises, 93–97

exercises, solutions for, 97–102

FROM clauses, overview of, 29–30

GROUP BY clause, overview of, 32–35

HAVING clause, overview of, 35–36

identifier names, delimiting of, 30

in temporal tables, 304–310

LIKE predicate, 71–73

logical query processing, overview, 27–29

nested and outer queries, defined, 133

NULLs, overview of, 54–58

of metadata, 89–92

OFFSET-FETCH filter, overview of, 46–47

operators and functions, overview of, 63–73

ORDER BY clause, overview of, 41–43

predicates and operators, overview of, 49–52

query clause, defined, 29

- query phrase, defined, 29
- scalar subqueries, 133
- SELECT clause, overview of, 36–41
- TOP filters, overview of, 44–46
- WHERE clause, overview of, 31–32
- window functions, overview of, 47–49

quoted identifiers, 62–63

R

RAND, 383

RANK, window function ranking, 216–219

RDBMSs (relational database management systems)

- ABC flavors (Appliance, Box, Cloud), 12–14
- defined, 1
- language independence, 2
- relational model, overview, 4–10

READ COMMITTED

- conflict detection, SNAPSHOT, 341–343
- default isolation levels, 323
- isolation, overview of, 332–337, 345
- SNAPSHOT and, 340, 343

READ COMMITTED LOCKS, 343

READ COMMITTED SNAPSHOT

- conflict detection, SNAPSHOT, 341–343
- default isolation levels, 323, 332–333
- isolation, overview, 339, 343–345

READ UNCOMMITTED, isolation overview, 332–334, 345

referenced table, foreign-key constraints, 23–24

referencing table, foreign-key constraints, 23–24

relation variable, use of term, 5

relational database management systems (RDBMSs)

- ABC flavors (Appliance, Box, Cloud), 12–14
- defined, 1
- language independence, 2
- relational model, overview, 4–10

relational model

- constraints, 7
- missing values, 6–7
- normalization, 7–10
- overview of, 4–10
- propositions, predicates, and relations, 5–6

REPEATABLE READ, 332–333, 335–337, 345

REPLACE, 66–67

REPLICATE, 67–68

resolution, batches and, 366

Resource database, SQL Server architecture, 16

resource types, locking of, 324–325

RESTART WITH, sequence objects and, 259

restoring tables. See temporal tables

RIGHT, 65

ROLLBACK TRAN (TRANSACTION)

- lock modes and compatibility, 322–324
- overview of, 319–322
- triggers and, 387

ROLLUP, 235–237

routines

- overview of, 383
- stored procedures, 385–386
- triggers, 387–390
- user-defined functions, 383–384

row constructors, 270

row versioning

- isolation and, 320–321
- READ COMMITTED SNAPSHOT, 343–345
- SNAPSHOT, 339–343

ROW_NUMBER

- EXCEPT ALL and, 200–201
- INTERSECT ALL and, 197–199
- window functions, overview of, 47–49
- window functions, ranking of, 216–219

rows, 4

ROWS BETWEEN

- offset window functions, 220–221
- window functions, overview of, 215

RTRIM, 69

running aggregates, subqueries and, 144–145

S

scalar subqueries

- defined, 133
- scalar variables and, 361–363
- self-contained subquery examples, 134–135

scalar UDFs (user-defined functions), 383–384

scalar variables, 361–363

SCHEMABINDING, views and, 176–177

schemas

- dbo-schema, 20
- identifier names, 30
- information schema views, metadata queries, 91
- object names, qualifying of, 29
- SQL Server architecture, 18–19

SCOPE_IDENTITY, 255–257

searched CASE expressions, 52–54

security

- database schemas and objects, SQL Server, 18–19
- stored procedures, 385

SELECT

SELECT

- INSERT SELECT, 251
 - pivoting with grouped queries, 226–227
 - variables and, 362–363
 - window functions, overview of, 215–216
 - window functions, ranking of, 218–219
- SELECT * FROM, views and, 172
- SELECT *, views and, 172
- SELECT INTO, 252–253
- SELECT statement. See also table expressions
 - all-at-once operations, 58–59
 - CASE expressions, overview of, 52–54
 - character data, working with, 60–63
 - date and time data, working with, 73–89
 - derived tables, assigning aliases, 163–164
 - elements of, 27–29
 - exercises, 93–97
 - exercises, solutions for, 97–102
 - FROM, overview, 29–30
 - GROUP BY, overview of, 32–35
 - HAVING, overview of, 35–36
 - LIKE predicate, 71–73
 - logical query processing order, 28
 - metadata, querying of, 89–92
 - NULLs, overview of, 54–58
 - OFFSET-FETCH filter, overview of, 46–47
 - operators and functions, overview of, 63–71
 - ORDER BY, overview of, 41–43
 - predicates and operators, overview of, 49–52
 - SELECT, overview of, 36–41
 - TOP filters, overview of, 44–46
 - views and, 172
 - WHERE, overview of, 31–32
 - window functions, overview of, 47–49
- self joins, 105–106
- self-contained subqueries
 - defined, 133
 - multivalued subquery examples, 136–139
 - scalar subquery examples, 134–135
- semicolon, use in statements, 29
- SEQUEL (Structured English QUery Language), 2
- sequence object, 258–262
- SERIALIZALBE
 - isolation, overview, 332–333, 337–338, 345
 - SNAPSHOT and, 340
- SERVERPROPERTY, 92
- session ID, troubleshooting blocks, 326–332
- SET
 - quoted identifiers, 62–63
 - UPDATE and, 267–268
 - variables and, 361–363
- set attributes
 - normalization forms, 7–10
 - propositions, predicates, and relations, 5–6
 - use of term, 5
- SET DEFAULT, 24
- SET NOCOUNT ON, 386
- SET NULL, 24
- set operators
 - EXCEPT, 199–201
 - exercises, 204–208
 - exercises, solutions for, 208–211
 - INTERSECT, 196–199
 - overview of, 193–194
 - precedence and, 201–202
 - UNION, 194–196
 - unsupported logical phrases, circumventing, 203–204
- set theory
 - overview of, 3
 - predicate logic and, 4
- simple CASE expressions, 52–54
- skipping option, OFFSET-FETCH, 46–47
- SMALLDATETIME
 - data types, 73–74
 - literals, 74–78
 - working with date and time separately, 78–79
- SNAPSHOT
 - conflict detection, 341–343
 - isolation, overview of, 339–343, 345
 - READ COMMITTED SNAPSHOT, 343–345
- snowflake dimension, 11
- solutions. See exercises
- sorting, ORDER BY clause and, 41–43
- sp_columns, 92
- sp_executesql stored procedure, 380–381
- sp_help, 91–92
- sp_helpconstraint, 92
- sp_tables, 91
- SQL (Structured Query Language), overview of, 1–3
 - database system types, 10–12
 - history and use of, 2–3
 - predicate logic, 4
 - relational model, overview of, 4–10
 - constraints, 7
 - missing values, 6–7
 - normalization, 7–10
 - propositions, predicates, and relations, 5–6
 - relational models, language independence, 2
 - set theory, overview of, 3

- standards for use, 2
 - statement termination, semicolons, 29
 - SQL Data Warehouse, 14
 - SQL Server
 - architecture overview, 12–19
 - ABC flavors (Appliance, Box, Cloud), 12–14
 - databases, 15–18
 - instances, 14–15
 - physical layout, 17–18
 - schemas and objects, 18–19
 - box product installation, 396–401
 - database engine installation, 396–401
 - file extensions, 18
 - operator precedence rules, 51–52
 - source code, downloading of, 402–404
 - SQL Server Books Online, 410–413
 - unique index, 22–23
 - SQL Server authenticated login, 17
 - SQL Server Management Studio (SSMS)
 - download and installation of, 402
 - getting started with, 404–410
 - session ID, troubleshooting blocks, 326–332
 - temporal tables, creating, 299
 - SQL_VARIANT, 262
 - star schema, overview of, 11–12
 - START WITH, sequence objects and, 258–259
 - storage, character data and, 60–61
 - stored procedures, 385–386
 - triggers, 387–390
 - string concatenation, 63–65
 - STRING_SPLIT, 71
 - Structured Query Language (SQL), overview of, 1–3
 - STUFF, 68
 - subqueries. *See also* table expressions
 - column names, substitution errors, 147–149
 - correlated subqueries
 - EXISTS and, 141–143
 - overview of, 139–141
 - exercises, 150–154
 - exercises, solutions for, 154–159
 - multivalued subqueries, defined, 133
 - NULL trouble, 145–147
 - overview of, 133
 - previous or next values, returning, 143–144
 - running aggregates and, 144–145
 - scalar subqueries, defined, 133
 - self-contained
 - defined, 133
 - multivalued subquery examples, 136–139
 - scalar subquery examples, 134–135
 - table subqueries, defined, 133
 - vs. joins, 137
 - window functions, overview of, 214
 - SUBSTRING, 65
 - subtraction (-), 51–52
 - SUM, 33–35. *See also* window functions
 - surrogate keys, identity property and, 254–257
 - SWITCHOFFSET, 83
 - syntax
 - batch as unit of parsing, 364
 - cross joins, 104–105
 - inner joins, 107–109
 - sys.dm_as_waiting_tasks, 330–331
 - sys.dm_exec_connections, 328
 - sys.dm_exec_input_buffer, 329
 - sys.dm_exec_requests, 330–331
 - sys.dm_exec_sessions, 329
 - sys.dm_exec_sql_text, 328
 - sys.dm_tran_locks, 328
 - SYSDATETIME
 - current date and time functions, 80–81
 - default constraints, defining, 25
 - INSERT SELECT and, 251
 - SYSDATETIMEOFFSET, 80–81
 - system versioned temporal tables. *See* temporal tables
 - SYSTEM_TIME, temporal table queries, 304–310
 - SYSTEM_VERSIONING, temporal table creation, 298–301
 - SYSUTCDATETIME, 80–81
- ## T
- table expressions
 - APPLY operator, 180–183
 - common table expressions (CTEs)
 - arguments, using, 167–168
 - column aliases, assigning, 167–168
 - defining multiple CTEs, 168–169
 - multiple references in, 169
 - overview, 167
 - recursive CTEs, 169–171
 - derived tables
 - arguments, use of, 164–165
 - column aliases, assigning, 163–164
 - multiple references, 166
 - nesting, 165–166
 - overview of, 161–162
 - exercises, 183–188, 287–291
 - exercises, solutions for, 188–192, 291–295

table subqueries

- table expressions *continued*
 - inline table-valued functions (TVFs), overview of, 171–172, 178–180
 - modifying data with, 276–278
 - overview of, 161
 - views
 - CHECK OPTION, 177–178
 - ENCRYPTION, 175–176
 - ORDER BY, 172–174
 - overview of, 171–172
 - SCHEMABINDING, 176–177
- table subqueries, defined, 133. See also table expressions
- tables
 - check constraints, defining, 24–25
 - creating tables, 19–22
 - defining data integrity, 22–25
 - foreign-key constraints, defining, 23–24
 - identifier names, 30
 - primary-key constraints, defining, 23
 - unique key constraints, defining, 23
- table-valued functions (TVFs), inline, 171–172
- table-valued parameters (TVPs), 378–379
- table-valued UDFs (user-defined functions), 383–384
- tempdb, SQL Server
 - architecture overview, 16–18
 - local temporary tables, 374–375
 - SNAPSHOT, 339–343
 - temporary table variables, 377–378
- temporal tables
 - creating, 298–301
 - exercises, 311–313
 - exercises, solutions for, 313–317
 - modifying data, 301–304
 - overview of, 297
 - querying data, 304–310
- temporary tables
 - global temporary tables, 376–377
 - local temporary tables, 374–375
 - overview of, 374
 - table types, 378–379
 - table variables, 377–378
- THEN, CASE expressions, 53–54
- tiebreakers, ORDER BY clause and, 45–46
- TIME
 - data types, 73–74
 - literals, 74–78
 - working with date and time separately, 78–79
- time data
 - aggregate window functions, 223–224
 - AT TIME ZONE, 84–85
 - CAST, CONVERT, and PARSE, 81–83
 - current date and time functions, 80–81
 - DATEADD, 85
 - DATEDIFF and DATEDIFF_BIG, 86–87
 - DATENAME, 88
 - DATEPART, 87
 - EOMONTH, 89
 - FROMPARTS, 88–89
 - ISDATE, 88
 - literals, 74–78
 - missing values, in outer joins, 116–118
 - SWITCHOFFSET, 83
 - types of, 73–74
 - working with date and time separately, 78–79
 - YEAR, MONTH, and DAY, 87–88
- TODATETIMEOFFSET, 83–84
- TOP (100) PERCENT, views and, 173–174
- TOP filters
 - data modification with, 278–280
 - derived tables and, 162
 - exercises, 287–291
 - exercises, solutions for, 291–295
 - overview of, 44–46
 - views and ORDER BY clause, 173–174
- transaction log, durability and, 321
- transactions
 - deadlocks, 345–348
 - exercises, 348–359
 - isolation
 - overview of, 332–333
 - READ COMMITTED, 334–337
 - READ COMMITTED SNAPSHOT, 343–345
 - READ UNCOMMITTED, 333–334
 - REPEATABLE READ, 335–337
 - SERIALIZABLE, 337–338
 - SNAPSHOT, 339–343
 - SNAPSHOT, conflict detection, 341–343
 - summary of, 345
 - locks and blocking, 322–325
 - locks and blocking, troubleshooting, 325–332
 - overview of, 319–322
 - vs. batches, 363
- triggers, 387–390
- troubleshooting
 - locks and blocking, 325–332
 - OUTPUT, use of, 280
 - table expressions, modifying data with, 276–278
- TRUE
 - check constraints and, 24–25
 - IF . . . ELSE flow element, 367–368

meaning of, 32
 NULLs, overview of, 54–58
 WHILE flow element, 368–370
 TRUNCATE, 262–266
 TRY . . . CATCH, error handling, 390–394
 TRY_, date and time functions, 81–83
 T-SQL
 coding style, 21–22
 database system types, 10–12
 language independence, 2
 predicate logic, 4
 relational model, overview of, 4–10
 constraints, 7
 missing values, 6–7
 normalization, 7–10
 propositions, predicates, and relations, 5–6
 set theory, 3–4
 SQL standards and, 3
 SQL, background information, 1–3
 whole, use of, 3
 tuples, 4–5
 type; propositions, predicates, and relations, 6

U

UDF (user-defined functions), routines, 383–384
 uncommitted reads, 333–334
 underscore (_) wildcard, 72
 Unicode data types, 50
 character data, 60–61
 UNION
 exercises, 204–208
 exercises, solutions for, 208–211
 precedence and, 201–202
 use of, 194–196
 UNION (DISTINCT), 194–196
 UNION ALL, 194–195
 unique constraint, data integrity, 23
 unique index, SQL Server
 primary-key constraints and, 22
 unique constraints, defining, 23
 UNIQUE, NULLs and, 58
 UNKNOWN
 check constraints and, 24–25
 IF . . . ELSE flow element, 367–368
 meaning of, 32
 NULLs, overview of, 54–58
 outer joins, filtering attributes, 118–119
 outer joins, in multi-join queries, 119–121
 scalar subqueries and, 135

 subqueries, NULL trouble and, 146–147
 WHILE flow element, 368–370
 unknown, missing values, 6–7
 UNPIVOT, 233–234
 unpivoting data, 230–234
 exercises, 241–244
 with APPLY, 231–233
 with UNPIVOT, 233–234
 UPDATE
 assignment UPDATE, 270–271
 based on a join, 268–270
 DML triggers (data manipulation events), 387–388
 exercises, 287–291
 exercises, solutions for, 291–295
 MERGE and, 274
 OUTPUT and, 283–284
 overview of, 266–268
 row-versioning isolation and, 339
 table expressions, modifying data with, 276–278
 views and, 172
 UPPER, 68–69
 USE, table creation, 20
 user-defined functions (UDF), routines, 383–384

V

VALUES, 250–251
 VAR element, 60–61
 VARCHAR, 21, 50, 60–61
 variables, as programmable objects
 batches and, 365
 overview of, 361–363
 temporary table variables, 377–378
 vector expressions, 270
 views, table expressions
 CHECK OPTION, 177–178
 ENCRYPTION, 175–176
 exercises, 183–188
 exercises, solutions for, 188–192
 ORDER BY and, 172–174
 overview of, 171–172
 SCHEMABINDING, 176–177
 virtual machine (VM), ABC flavors (Appliance, Box, Cloud), 13–14

W

WHEN MATCHED THEN, 273–275
 WHEN NOT MATCHED THEN, 273–275

WHEN, CASE expressions

WHEN, CASE expressions, 53–54

WHERE

CASE expressions, overview of, 52–54

derived tables, arguments, 164–165

derived tables, assigning aliases, 163

logical query processing order, 28

outer joins, filtering attributes, 118–119

outer joins, fundamentals of, 115–116

overview of, 31–32

UPDATE and, 267–268

UPDATE based on a join, 268–270

WHILE flow element, 368–370

whole, defined, 3

wildcards

% (percent) wildcard, 71–72

_ (underscore) wildcard, 72

[^character list or range] wildcard, 73

[character-character] wildcard, 72–73

[list of characters] wildcard, 72

LIKE predicate and, 71–73

window functions

aggregate window functions, 221–224

exercises, 241–244

exercises, solutions for, 245–248

offset functions, 219–221

overview of, 47–49, 213–216

ranking of, 216–219

Windows authenticated login, SQL Server databases, 17

WITH NOCHECK, 25

WITH TIES, ORDER BY clause and, 45–46

WITH, common table expressions (CTEs), 167

X

XACT_ABORT, 320

Y

YEAR, 87–88

filtering date ranges, 79–80

YTD, aggregate window functions, 223–224