

LEARNING MADE EASY



4th Edition

SQL

ALL-IN-ONE

for
dummies[®]
A Wiley Brand



Allen G. Taylor
with Richard Blum



SQL

ALL-IN-ONE

4th Edition

by Allen G. Taylor
with Richard Blum

for
dummies[®]
A Wiley Brand

SQL All-in-One For Dummies®, 4th Edition

Published by: **John Wiley & Sons, Inc.**, 111 River Street, Hoboken, NJ 07030-5774, www.wiley.com

Copyright © 2024 by John Wiley & Sons, Inc., Hoboken, New Jersey

Media and software compilation copyright © 2024 by John Wiley & Sons, Inc. All rights reserved.

Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For technical support, please visit <https://hub.wiley.com/community/support/dummies>.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2024933440

ISBN 978-1-394-24229-0 (pbk); ISBN 978-1-394-24232-0 (ebk); ISBN 978-1-394-24231-3 (ebk)

Contents at a Glance

Introduction	1
Book 1: Getting Started with SQL	5
CHAPTER 1: Understanding Relational Databases	7
CHAPTER 2: Modeling a System	29
CHAPTER 3: Getting to Know SQL	53
CHAPTER 4: SQL and the Relational Model	65
CHAPTER 5: Knowing the Major Components of SQL	75
CHAPTER 6: Drilling Down to the SQL Nitty-Gritty	99
Book 2: Developing Relational Databases	133
CHAPTER 1: System Development Overview	135
CHAPTER 2: Building a Database Model	151
CHAPTER 3: Balancing Performance and Correctness	169
CHAPTER 4: Creating a Database with SQL	203
Book 3: Writing SQL Queries	217
CHAPTER 1: Values, Variables, Functions, and Expressions	219
CHAPTER 2: SELECT Statements and Modifying Clauses	249
CHAPTER 3: Querying Multiple Tables with Subqueries	295
CHAPTER 4: Querying Multiple Tables with Relational Operators	323
CHAPTER 5: Cursors	345
Book 4: Securing Your Data	357
CHAPTER 1: Protecting Against Hardware Failure and External Threats	359
CHAPTER 2: Protecting Against User Errors and Conflicts	391
CHAPTER 3: Assigning Access Privileges	419
CHAPTER 4: Error Handling	431
Book 5: Programming with SQL	447
CHAPTER 1: Database Development Environments	449
CHAPTER 2: Interfacing SQL to a Procedural Language	455
CHAPTER 3: Using SQL in an Application Program	461
CHAPTER 4: Designing a Sample Application	477
CHAPTER 5: Building an Application	497
CHAPTER 6: Understanding SQL's Procedural Capabilities	513
CHAPTER 7: Connecting SQL to a Remote Database	531

Book 6: Working with Advanced Data Types in SQL: XML, JSON, and PGQ	545
CHAPTER 1: Using XML with SQL	547
CHAPTER 2: Storing XML Data in SQL Tables	575
CHAPTER 3: Retrieving Data from XML Documents	599
CHAPTER 4: Using JSON with SQL	617
CHAPTER 5: Exploring Property Graph Queries	633
Book 7: Optimizing Your Database	645
CHAPTER 1: Tuning the Database	647
CHAPTER 2: Tuning the Environment	659
CHAPTER 3: Finding and Eliminating Performance Bottlenecks	681
Book 8: Appendixes	709
CHAPTER 1: SQL:2023 Reserved Words	711
CHAPTER 2: Glossary	721
Index	731

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book	3
Beyond the Book	3
Where to Go from Here	4
BOOK 1: GETTING STARTED WITH SQL	5
CHAPTER 1: Understanding Relational Databases	7
Understanding Why Today's Databases Are Better than Early Databases	8
Irreducible complexity	8
Managing data with complicated programs	8
Managing data with simple programs	11
Which type of organization is better?	11
Databases, Queries, and Database Applications	12
Making data useful	12
Retrieving the data you want — and only the data you want. . . .	12
Examining Competing Database Models	14
Looking at the historical background of the competing models . . .	14
The hierarchical database model	15
The network database model	19
The relational database model	21
The object-oriented database model	25
The object-relational database model	25
The nonrelational NoSQL model	25
Why the Relational Model Won	26
CHAPTER 2: Modeling a System	29
Capturing the Users' Data Model	29
Identifying and interviewing stakeholders	30
Reconciling conflicting requirements	31
Obtaining stakeholder buy-in	31
Translating the Users' Data Model to a Formal Entity-Relationship Model	32
Entity-Relationship modeling techniques	33
Drawing Entity-Relationship diagrams	38
Understanding advanced ER model concepts	41
A simple example of an ER model	45
A slightly more complex example	47

	Problems with complex relationships	50
	Simplifying relationships using normalization	51
	Translating an ER model into a relational model	51
CHAPTER 3:	Getting to Know SQL	53
	Where SQL Came From	53
	Knowing What SQL Does	54
	The ISO/IEC SQL Standard	55
	Knowing What SQL Does Not Do	55
	Choosing and Using an Available DBMS Implementation	56
	Microsoft Access	57
	Microsoft SQL Server	61
	IBM DB2	61
	Oracle Database	62
	Sybase SQL Anywhere	62
	MySQL	62
	PostgreSQL	63
CHAPTER 4:	SQL and the Relational Model	65
	Sets, Relations, Multisets, and Tables	66
	Functional Dependencies	67
	Keys	68
	Views	69
	Users	70
	Privileges	70
	Schemas	71
	Catalogs	72
	Connections, Sessions, and Transactions	72
	Routines	73
	Paths	73
CHAPTER 5:	Knowing the Major Components of SQL	75
	Creating a Database with the Data Definition Language	75
	The containment hierarchy	76
	Creating tables	77
	Specifying columns	77
	Creating other objects	78
	Modifying tables	85
	Removing tables and other objects	85
	Operating on Data with the Data Manipulation Language (DML)	86
	Retrieving data from a database	86
	Adding data to a table	87
	Updating data in a table	90
	Deleting data from a table	93

	Updating views doesn't make sense.	94
	Maintaining Security in the Data Control Language (DCL)	95
	Granting access privileges	95
	Revoking access privileges.	96
	Preserving database integrity with transactions.	96
CHAPTER 6:	Drilling Down to the SQL Nitty-Gritty	99
	Executing SQL Statements	99
	Interactive SQL	100
	Challenges to combining SQL with a host language	101
	Embedded SQL	101
	Module language.	104
	Using Reserved Words Correctly.	105
	SQL's Data Types	105
	Exact numerics.	106
	Approximate numerics.	108
	Character strings	110
	Binary strings	112
	Booleans	113
	Datetimes	113
	Intervals	115
	XML type	116
	ROW type	116
	Collection types	118
	REF types.	119
	JSON types	119
	User-defined types	120
	Data type summary.	124
	Handling Null Values.	125
	Applying Constraints.	126
	Column constraints.	126
	Table constraints	128
	Foreign key constraints.	129
	Assertions	131
	BOOK 2: DEVELOPING RELATIONAL DATABASES	133
CHAPTER 1:	System Development Overview	135
	The Components of a Database System	136
	The database	136
	The database engine.	137
	The DBMS front end	137
	The database application	137
	The user	138

	The System Development Life Cycle	138
	Definition phase	139
	Requirements phase	140
	Evaluation phase	142
	Design phase	144
	Implementation phase	147
	Final Documentation and Testing phase	148
	Maintenance phase	150
CHAPTER 2:	Building a Database Model	151
	Finding and Listening to Interested Parties	152
	Your immediate supervisor	152
	The users	153
	The standards organization	153
	Upper management	154
	Building Consensus	154
	Gauging what people want	156
	Arriving at a consensus	156
	Building a Relational Model	156
	Reviewing the three database traditions	157
	Knowing what a relation is	158
	Functional dependencies	158
	Keys	159
	Being Aware of the Danger of Anomalies	160
	Eliminating anomalies	161
	Examining the higher normal forms	164
	The Database Integrity versus Performance Tradeoff	166
CHAPTER 3:	Balancing Performance and Correctness	169
	Designing a Sample Database	170
	The ER model for Honest Abe's	170
	Converting an ER model into a relational model	172
	Normalizing a relational model	172
	Handling binary relationships	174
	A sample conversion	179
	Maintaining Integrity	182
	Entity integrity	182
	Domain integrity	183
	Referential integrity	184
	Avoiding Data Corruption	186
	Speeding Data Retrievals	187
	Hierarchical storage	188
	Full table scans	189

Working with Indexes	189
Creating the right indexes	190
Indexes and the ANSI/ISO standard	190
Index costs	191
Query type dictates the best index.	191
Data structures used for indexes	193
Indexes, sparse and dense.	194
Index clustering	195
Composite indexes	195
Index effect on join performance	196
Table size as an indexing consideration.	196
Indexes versus full table scans	197
Reading SQL Server Execution Plans	197
Robust execution plans	197
A sample database	198
CHAPTER 4: Creating a Database with SQL.	203
First Things First: Planning Your Database	203
Building Tables.	204
Locating table rows with keys	204
Using the CREATE TABLE statement	207
Setting Constraints	209
Column constraints.	209
Table constraints.	210
Keys and Indexes.	210
Ensuring Data Validity with Domains	210
Establishing Relationships between Tables.	211
Altering Table Structure	215
Deleting Tables	215
BOOK 3: WRITING SQL QUERIES.	217
CHAPTER 1: Values, Variables, Functions, and Expressions	219
Entering Data Values.	219
Row values have multiple parts.	220
Identifying values in a column.	220
Literal values don't change.	220
Variables vary.	222
Special variables hold specific values.	222
Working with Functions	223
Summarizing data with set functions	223
Dissecting data with value functions	227
Using Expressions	237
Numeric value expressions	237
String value expressions.	238

	Datetime value expressions	238
	Interval value expressions	239
	Boolean value expressions	240
	Array value expressions	241
	Conditional value expressions	241
	Converting data types with a CAST expression	244
	Row value expressions	246
CHAPTER 2:	SELECT Statements and Modifying Clauses	249
	Finding Needles in Haystacks with the SELECT Statement	249
	Modifying Clauses	250
	FROM clauses	250
	WHERE clauses	251
	GROUP BY clauses	270
	HAVING clauses	273
	ORDER BY clauses	274
	Tuning Queries	276
	SELECT DISTINCT	277
	Temporary tables	280
	The ORDER BY clause	285
	The HAVING clause	289
	The OR logical connective	293
CHAPTER 3:	Querying Multiple Tables with Subqueries	295
	What Is a Subquery?	295
	What Subqueries Do	296
	Subqueries that return multiple values	296
	Subqueries that return a single value	298
	Quantified subqueries return a single value	301
	Correlated subqueries	304
	Using Subqueries in INSERT, DELETE, and UPDATE Statements	309
	Tuning Considerations for Statements Containing Nested Queries	312
	Tuning Correlated Subqueries	318
CHAPTER 4:	Querying Multiple Tables with Relational Operators	323
	UNION	324
	UNION ALL	326
	UNION CORRESPONDING	326
	INTERSECT	327
	EXCEPT	329
	JOINS	330

Cartesian product or cross join	330
Equi-join	332
Natural join.	334
Condition join.	335
Column-name join	335
Inner join.	337
Outer join	337
ON versus WHERE	341
Join Conditions and Clustering Indexes	342
CHAPTER 5: Cursors	345
Declaring a Cursor.	346
The query expression	347
Ordering the query result set	347
Updating table rows	349
Sensitive versus insensitive cursors	349
Scrolling a cursor	351
Holding a cursor	351
Declaring a result set cursor	351
Opening a Cursor	352
Operating on a Single Row.	353
FETCH syntax	354
Absolute versus relative fetches	355
Deleting a row	355
Updating a row	355
Closing a Cursor.	356
BOOK 4: SECURING YOUR DATA	357
CHAPTER 1: Protecting Against Hardware Failure and External Threats	359
What Could Possibly Go Wrong?.	360
Equipment failure	360
Platform instability	362
Database design flaws	363
Data-entry errors.	363
Operator error	363
Taking Advantage of RAID	363
Striping	364
RAID levels	366
Backing Up Your System.	368
Preparation for the worst.	368
Full or incremental backup	369
Frequency.	369
Backup maintenance	369

Coping with Internet Threats	370
Viruses	370
Trojan horses	372
Worms	373
Denial-of-service attacks	374
Ransomware	374
SQL injection attacks	374
Phishing scams	387
Zombie spambots	387
Installing Layers of Protection	388
Network-layer firewalls	388
Application-layer firewalls	388
Antivirus software	388
Vulnerabilities, exploits, and patches	388
Education	389
Alertness	389
CHAPTER 2: Protecting Against User Errors and Conflicts	391
Reducing Data-Entry Errors	392
Data types: The first line of defense	392
Constraints: The second line of defense	392
Sharp-eyed humans: The third line of defense	393
Coping with Errors in Database Design	393
Handling Programming Errors	394
Solving Concurrent-Operation Conflicts	394
Passing the ACID Test: Atomicity, Consistency, Isolation, and Durability	396
Operating with Transactions	397
Using the SET TRANSACTION statement	397
Starting a transaction	398
Committing a transaction	401
Rolling back a transaction	401
Implementing deferrable constraints	404
Getting Familiar with Locking	409
Two-phase locking	409
Granularity	410
Deadlock	410
Tuning Locks	411
Measuring performance with throughput	412
Eliminating unneeded locks	412
Shortening transactions	412
Weakening isolation levels (ver-r-ry carefully)	413
Controlling lock granularity	414
Scheduling DDL statements correctly	414

	Partitioning insertions	414
	Cooling hot spots.	415
	Tuning the deadlock interval.	415
	Enforcing Serializability with Timestamps.	415
	Tuning the Recovery System	418
CHAPTER 3:	Assigning Access Privileges	419
	Working with the SQL Data Control Language	419
	Identifying Authorized Users	420
	Understanding user identifiers	420
	Getting familiar with roles	420
	Classifying Users	422
	Granting Privileges	422
	Looking at data	423
	Deleting data	424
	Adding data	424
	Changing data	424
	Referencing data in another table	425
	Using certain database facilities	426
	Responding to an event	427
	Defining new data types.	427
	Executing an SQL statement	427
	Doing it all.	427
	Passing on the power	428
	Revoking Privileges	428
	Granting Roles	430
	Revoking Roles.	430
CHAPTER 4:	Error Handling	431
	Identifying Error Conditions.	432
	Getting to Know SQLSTATE	432
	Handling Conditions	434
	Handler declarations.	435
	Handler actions and handler effects	435
	Conditions that aren't handled	437
	Dealing with Execution Exceptions: The WHENEVER Clause.	437
	Getting More Information: The Diagnostics Area.	438
	The diagnostics header area	439
	The diagnostics detail area	440
	Examining an Example Constraint Violation	442
	Adding Constraints to an Existing Table	444
	Interpreting SQLSTATE Information	444
	Handling Exceptions	445

BOOK 5: PROGRAMMING WITH SQL	447
CHAPTER 1: Database Development Environments	449
Microsoft Access	449
The Jet engine.	450
DAO	450
ADO	450
ODBC	451
OLE DB	451
Files with the .mdb extension	451
The Access Database Engine	451
Microsoft SQL Server	451
IBM Db2	452
Oracle 23c	453
SQL Anywhere	453
PostgreSQL	453
MySQL	454
CHAPTER 2: Interfacing SQL to a Procedural Language	455
Building an Application with SQL and a Procedural Language	455
Access and VBA	456
SQL Server and the .NET languages	457
MySQL and C++.NET or C#	458
MySQL and C	458
MySQL and Perl	459
MySQL and Python	459
MySQL and PHP	459
MySQL and Java	459
Oracle SQL and Java	459
Db2 and Java	460
CHAPTER 3: Using SQL in an Application Program	461
Comparing SQL with Procedural Languages	462
Classic procedural languages	462
Object-oriented procedural languages	463
Nonprocedural languages	463
Difficulties in Combining SQL with a Procedural Language	464
Challenges of using SQL with a classical procedural language	464
Challenges of using SQL with an object-oriented procedural language	465
Embedding SQL in an Application	466
Embedding SQL in an Oracle Pro*C application	467
Embedding SQL in a Java application	470
Using SQL in a Perl application	470
Embedding SQL in a PHP application	471

	Using SQL with a Visual Basic .NET application	471
	Using SQL with other .NET languages	472
	Using SQL Modules with an Application	472
	Module declarations	473
	Module procedures	474
	Modules in Oracle	474
CHAPTER 4:	Designing a Sample Application	477
	Understanding the Client's Problem	478
	Approaching the Problem	478
	Interviewing the stakeholders	478
	Drafting a detailed statement of requirements	479
	Following up with a proposal	479
	Determining the Deliverables	480
	Finding out what's needed now and later	481
	Planning for organization growth	481
	Nailing down project scope	482
	Building an Entity-Relationship Model	483
	Determining what the entities are	484
	Relating the entities to one another	484
	Transforming the Model	487
	Eliminating any many-to-many relationships	487
	Normalizing the ER model	490
	Creating Tables	491
	Changing Table Structure	495
	Removing Tables	495
	Designing the User Interface	496
CHAPTER 5:	Building an Application	497
	Designing from the Top Down	497
	Determining what the application should include	498
	Designing the user interface	498
	Connecting the user interface to the database	499
	Coding from the Bottom Up	501
	Preparing to build the application	501
	Creating the application's building blocks	509
	Gluing everything together	510
	Testing, Testing, Testing	510
	Fixing the bugs	511
	Turning naive users loose	511
	Bringing on the hackers	511
	Fixing the newly found bugs	511
	Retesting everything one last time	512

CHAPTER 6: Understanding SQL's Procedural Capabilities	513
Embedding SQL Statements in Your Code	514
Introducing Compound Statements	514
Atomicity	515
Variables	516
Cursors	516
Assignment	517
Following the Flow of Control Statements	517
IF . . . THEN . . . ELSE . . . END IF	518
CASE . . . END CASE	518
LOOP . . . END LOOP	520
LEAVE	520
WHILE . . . DO . . . END WHILE	521
REPEAT . . . UNTIL . . . END REPEAT	521
FOR . . . DO . . . END FOR	522
ITERATE	522
Using Stored Procedures	523
Working with Triggers	524
Trigger events	525
Trigger action time	526
Triggered actions	526
Triggered SQL statement	526
Using Stored Functions	527
Passing Out Privileges	528
Using Stored Modules	528
CHAPTER 7: Connecting SQL to a Remote Database	531
Native Drivers	532
ODBC and Its Major Components	533
Application	534
Driver manager	535
Drivers	535
Data sources	537
What Happens When the Application Makes a Request	537
Using handles to identify objects	538
Following the six stages of an ODBC operation	538
BOOK 6: WORKING WITH ADVANCED DATA TYPES IN SQL: XML, JSON, AND PGQ	545
CHAPTER 1: Using XML with SQL	547
Introducing XML	548
Knowing the Parts of an XML Document	549

XML declaration	549
Elements	550
Attributes	551
Entity references	552
Numeric character references	553
Using XML Schema	553
Relating SQL to XML	554
Using the XML Data Type	555
When to use the XML type	555
When not to use the XML type	556
Mapping SQL to XML	557
Mapping character sets to XML	557
Mapping identifiers to XML	557
Mapping data types to XML	558
Mapping nonpredefined data types to XML	559
Mapping tables to XML	564
Handling null values	564
Creating an XML schema for an SQL table	565
Operating on XML Data with SQL Functions	566
XMLELEMENT	567
XMLFOREST	567
XMLCONCAT	568
XMLAGG	568
XMLCOMMENT	569
XMLPARSE	569
XMLPI	570
XMLQUERY	570
XMLCAST	571
Working with XML Predicates	571
DOCUMENT	571
CONTENT	572
XMLEXISTS	572
VALID	572
CHAPTER 2: Storing XML Data in SQL Tables	575
Inserting XML Data into an SQL Pseudotable	575
Creating a Table to Hold XML Data	577
Updating XML Documents	578
Discovering Oracle's Tools for Updating XML Data in a Table	579
APPENDCHILDXML	579
INSERTCHILDXML	580
INSERTXMLBEFORE	581
DELETXML	582
UPDATERXML	583

	Introducing Microsoft's Tools for Updating XML Data in a Table	584
	Inserting data into a table using OPENXML	584
	Using updategrams to map data into database tables	585
	Using an updategram namespace and keywords	585
	Specifying a mapping schema	587
CHAPTER 3:	Retrieving Data from XML Documents	599
	XQuery	600
	Where XQuery came from	600
	What XQuery requires	601
	XQuery functionality	601
	Usage scenarios	602
	FLWOR Expressions	607
	The for clause	608
	The let clause	609
	The where clause	610
	The order by clause	611
	The return clause	611
	XQuery versus SQL	612
	Comparing XQuery's FLWOR expression with SQL's SELECT expression	613
	Relating XQuery data types to SQL data types	613
CHAPTER 4:	Using JSON with SQL	617
	Using JSON with SQL	617
	The SQL/JSON Data Model	618
	SQL/JSON items	618
	SQL/JSON sequences	619
	Parsing JSON	620
	Serializing JSON	620
	SQL/JSON Functions	620
	Query functions	620
	Constructor functions	626
	IS JSON predicate	628
	JSON nulls and SQL nulls	629
	SQL/JSON Path Language	629
	SQL:2023 JSON Enhancements	629
	The JSON data type	630
	Additional functions for the JSON data type	631
CHAPTER 5:	Exploring Property Graph Queries	633
	What Are Property Graph Queries?	633
	Looking at node and edge properties	634
	Connecting nodes by multiple edges	634

Using edges to connect a node to itself	635
Following paths with SQL queries	635
Examining SQL/PgQ	635
Working with SQL/PgQ	637
Building the property graph tables	637
Adding data to node and edge tables	640
Querying data in graph tables	642
BOOK 7: OPTIMIZING YOUR DATABASE	645
CHAPTER 1: Tuning the Database	647
Analyzing the Workload	648
Considering the Physical Design	649
Choosing the Right Indexes	650
Avoiding unnecessary indexes	650
Choosing a column to index	651
Using multicolumn indexes	652
Clustering indexes	652
Choosing an index type	654
Weighing the cost of index maintenance	654
Using composite indexes	654
Tuning Indexes	655
Tuning Queries	656
Tuning Transactions	657
Separating User Interactions from Transactions	657
Minimizing Traffic between Application and Server	658
Precompiling Frequently Used Queries	658
CHAPTER 2: Tuning the Environment	659
Surviving Failures with Minimum Data Loss	660
What happens to transactions when no failure occurs?	660
What happens when a failure occurs and a transaction is still active?	661
Tuning the Recovery System	661
Volatile and nonvolatile memory	662
Memory system hierarchy	663
Putting logs and transactions on different disks	664
Tuning write operations	667
Performing database dumps	668
Setting checkpoints	669
Optimizing batch transactions	670
Tuning the Operating System	670
Scheduling threads	670
Determining database buffer size	674
Tuning the page usage factor	675

Maximizing the Hardware You Have	675
Optimizing the placement of code and data on hard disks	676
Tuning the page replacement algorithm	676
Tuning the disk controller cache	677
Adding Hardware	677
Faster processor	678
More RAM	678
Faster hard disks	678
More hard disks	679
Solid State Disk (SSD)	679
RAID arrays	679
Working in Multiprocessor Environments	680

CHAPTER 3: Finding and Eliminating Performance

Bottlenecks	681
Pinpointing the Problem	682
Slow query	682
Slow update	682
Determining the Possible Causes of Trouble	683
Problems with indexes	683
Pitfalls in communication	685
Determining whether hardware is robust enough and configured properly	687
Implementing General Principles: A First Step Toward Improving Performance	687
Avoid direct user interaction	687
Examine the application/database interaction	688
Don't ask for columns that you don't need	688
Don't use cursors unless you absolutely have to	689
Precompiled queries	689
Tracking Down Bottlenecks	689
Isolating performance problems	690
Performing a top-down analysis	690
Partitioning	692
Locating hotspots	693
Analyzing Query Efficiency	693
Using query analyzers	693
Finding problem queries	701
Managing Resources Wisely	705
The disk subsystem	705
The database buffer manager	706
The logging subsystem	707
The locking subsystem	707

BOOK 8: APPENDIXES	709
CHAPTER 1: SQL:2023 Reserved Words	711
CHAPTER 2: Glossary	721
INDEX	731

Introduction

SQL is the internationally recognized standard language for dealing with data in relational databases. Developed by IBM, SQL became an international standard in 1986. The standard was updated in 1989, 1992, 1999, 2003, 2008, 2011, 2016, and 2023. It continues to evolve and gain capability. Database vendors continually update their products to incorporate the new features of the ISO/IEC standard. (For the curious out there, ISO is the International Organization for Standardization, and IEC is the International Electrotechnical Commission.)

SQL isn't a general-purpose language, such as C++ or Java. Instead, it's strictly designed to deal with data in relational databases. With SQL, you can carry out all the following tasks:

- » Create a database, including all tables and relationships.
- » Fill database tables with data.
- » Change the data in database tables.
- » Delete data from database tables.
- » Retrieve specific information from database tables.
- » Grant and revoke access to database tables.
- » Protect database tables from corruption due to access conflicts or user mistakes.

About This Book

This book isn't just about SQL; it's also about how SQL fits into the process of creating and maintaining databases and database applications. In this book, I cover how SQL fits into the larger world of application development and how it handles data coming in from other computers, which may be on the other side of the world or even in interplanetary space.

Here are some of the things you can do with this book:

- » Create a model of a proposed system and then translate that model into a database.
- » Find out about the capabilities and limitations of SQL.
- » Discover how to develop reliable and maintainable database systems.
- » Create databases.
- » Speed database queries.
- » Protect databases from hardware failures, software bugs, and Internet attacks.
- » Control access to sensitive information.
- » Write effective database applications.
- » Deal with data from a variety of nontraditional data sources by using XML.

I've structured this book modularly — that is, it's designed so that you can easily find just the information you need — so you don't have to read whatever doesn't pertain to your task at hand. Here and there throughout the book, I include sidebars containing interesting information that isn't necessarily integral to the discussion at hand; feel free to skip them. You also don't have to read text marked with the Technical Stuff icons, which parses out über-techy tidbits (which may or may not be your cup of tea).

Within this book, you may note that some web addresses break across two lines of text. If you're reading this book in print and want to visit one of these web pages, simply key in the web address exactly as it's noted in the text, pretending as though the line break doesn't exist. If you're reading this as an e-book, you've got it easy — just click the web address to be taken directly to the web page.

Foolish Assumptions

I know that this is a *For Dummies* book, but I don't really expect that you're a dummy. In fact, I assume that you're a very smart person. After all, you decided to read this book, which is a sign of high intelligence indeed. Therefore, I assume that you may want to do a few things, such as re-create some of the examples in the book. You may even want to enter some SQL code and execute it. To do that, you need at the very least an SQL editor and more likely also a database management system (DBMS) of some sort. Many choices are available, both proprietary

and open source. I mention several of these products at various places throughout the book but don't recommend any one in particular. Any product that complies with the ISO/IEC international SQL standard should be fine.

Take claims of ISO/IEC compliance with a grain of salt, however. No DBMS available today is 100 percent compliant with the ISO/IEC SQL standard. For that reason, some of the code examples I give in this book may not work in the particular SQL implementation that you're using. The code samples I use in this book are consistent with the international standard rather than with the syntax of any particular implementation unless I specifically state that the code is for a particular implementation.

Icons Used in This Book

For *Dummies* books are known for those helpful icons that point you in the direction of really great information. This section briefly describes the icons used in this book.



TIP

The Tip icon points out helpful information that's likely to make your job easier.



REMEMBER

This icon marks a generally interesting and useful fact — something that you may want to remember for later use.



WARNING

The Warning icon highlights lurking danger. When you see this icon, pay attention, and proceed with caution.



TECHNICAL
STUFF

This icon denotes techie stuff nearby. If you're not feeling very techie, you can skip this info.

Beyond the Book

In addition to what you're reading right now, this book comes with a free access-anywhere Cheat Sheet that includes information on SQL system development, normalizing data, and SQL data types and functions. To get this Cheat Sheet, simply go to www.dummies.com and type **SQL All-in-One For Dummies Cheat Sheet** in the Search box.

Where to Go from Here

Book 1 is the place to go if you're just getting started with databases. It explains why databases are useful and describes the different types. It focuses on the relational model and describes SQL's structure and features.

Book 2 goes into detail on how to build a database that's reliable as well as responsive. Unreliable databases are much too easy to create, and this minibook tells you how to avoid the pitfalls that lie in wait for the unwary.

Go directly to Book 3 if your database already exists and you just want to know how to use SQL to pull from it the information you want.

Book 4 is primarily aimed at the database administrator (DBA) rather than the database application developer or user. It discusses how to build a robust database system that resists data corruption and data loss.

Book 5 is for the application developer. In addition to discussing how to write a database application, it gives an example that describes in a step-by-step manner how to build a reliable application.

If you're already an old hand at SQL and just want to know how to handle data in XML or JSON format in your SQL database, or if you'd like to dive into the property graph database world, Book 6 is for you.

Book 7 gives you a wide variety of techniques for improving the performance of your database. This minibook is the place to go if your database is operating — but not as well as you think it should. Most of these techniques are things that the DBA can do, rather than the application developer or the database user. If your database isn't performing the way you think it should, take it up with your DBA. She can do a few things that could help immensely.

Book 8 is a handy reference that helps you quickly find the meaning of a word you've encountered or see why an SQL statement that you entered didn't work as expected. (Maybe you used a reserved word without realizing it.)

1

Getting Started with SQL

Contents at a Glance

CHAPTER 1: Understanding Relational Databases	7
Understanding Why Today's Databases Are Better than Early Databases	8
Databases, Queries, and Database Applications	12
Examining Competing Database Models	14
Why the Relational Model Won	26
CHAPTER 2: Modeling a System	29
Capturing the Users' Data Model	29
Translating the Users' Data Model to a Formal Entity-Relationship Model	32
CHAPTER 3: Getting to Know SQL	53
Where SQL Came From	53
Knowing What SQL Does	54
The ISO/IEC SQL Standard	55
Knowing What SQL Does Not Do	55
Choosing and Using an Available DBMS Implementation	56
CHAPTER 4: SQL and the Relational Model	65
Sets, Relations, Multisets, and Tables	66
Functional Dependencies	67
Keys	68
Views	69
Users	70
Privileges	70
Schemas	71
Catalogs	72
Connections, Sessions, and Transactions	72
Routines	73
Paths	73
CHAPTER 5: Knowing the Major Components of SQL	75
Creating a Database with the Data Definition Language	75
Operating on Data with the Data Manipulation Language (DML) ..	86
Maintaining Security in the Data Control Language (DCL)	95
CHAPTER 6: Drilling Down to the SQL Nitty-Gritty	99
Executing SQL Statements	99
Using Reserved Words Correctly	105
SQL's Data Types	105
Handling Null Values	125
Applying Constraints	126