

SECOND EDITION

SQL Cookbook

*Query Solutions and Techniques
for All SQL Users*

Anthony Molinaro and Robert de Graaf

Beijing • Boston • Farnham • Sebastopol • Tokyo

O'REILLY®

SQL Cookbook

by Anthony Molinaro and Robert de Graaf

Copyright © 2021 Robert de Graaf. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<http://oreilly.com>). For more information, contact our corporate/institutional sales department: 800-998-9938 or corporate@oreilly.com.

Acquisitions Editor: Jessica Haberman

Development Editor: Virginia Wilson

Production Editor: Kate Galloway

Copyeditor: Kim Wimpsett

Proofreader: nSight, Inc.

Indexer: WordCo Indexing Services, Inc.

Interior Designer: David Futato

Cover Designer: Karen Montgomery

Illustrator: O'Reilly Media

December 2005: First Edition

December 2020: Second Edition

Revision History for the Second Edition

2020-11-03: First Release

See <http://oreilly.com/catalog/errata.csp?isbn=9781492077442> for release details.

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *SQL Cookbook*, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

The views expressed in this work are those of the authors, and do not represent the publisher's views. While the publisher and the authors have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the authors disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

This work is part of a collaboration between O'Reilly and Yugabyte. See our [statement of editorial independence](#).

978-1-492-07744-2

[LSI]

Table of Contents

Preface.....	xi
1. Retrieving Records.....	1
1.1 Retrieving All Rows and Columns from a Table	1
1.2 Retrieving a Subset of Rows from a Table	2
1.3 Finding Rows That Satisfy Multiple Conditions	2
1.4 Retrieving a Subset of Columns from a Table	3
1.5 Providing Meaningful Names for Columns	4
1.6 Referencing an Aliased Column in the WHERE Clause	5
1.7 Concatenating Column Values	6
1.8 Using Conditional Logic in a SELECT Statement	7
1.9 Limiting the Number of Rows Returned	8
1.10 Returning n Random Records from a Table	10
1.11 Finding Null Values	11
1.12 Transforming Nulls into Real Values	12
1.13 Searching for Patterns	13
1.14 Summing Up	14
2. Sorting Query Results.....	15
2.1 Returning Query Results in a Specified Order	15
2.2 Sorting by Multiple Fields	16
2.3 Sorting by Substrings	17
2.4 Sorting Mixed Alphanumeric Data	18
2.5 Dealing with Nulls When Sorting	21
2.6 Sorting on a Data-Dependent Key	27
2.7 Summing Up	28

3. Working with Multiple Tables.....	29
3.1 Stacking One Rowset atop Another	29
3.2 Combining Related Rows	31
3.3 Finding Rows in Common Between Two Tables	33
3.4 Retrieving Values from One Table That Do Not Exist in Another	34
3.5 Retrieving Rows from One Table That Do Not Correspond to Rows in Another	40
3.6 Adding Joins to a Query Without Interfering with Other Joins	42
3.7 Determining Whether Two Tables Have the Same Data	44
3.8 Identifying and Avoiding Cartesian Products	51
3.9 Performing Joins When Using Aggregates	52
3.10 Performing Outer Joins When Using Aggregates	57
3.11 Returning Missing Data from Multiple Tables	60
3.12 Using NULLs in Operations and Comparisons	64
3.13 Summing Up	65
4. Inserting, Updating, and Deleting.....	67
4.1 Inserting a New Record	68
4.2 Inserting Default Values	68
4.3 Overriding a Default Value with NULL	70
4.4 Copying Rows from One Table into Another	70
4.5 Copying a Table Definition	71
4.6 Inserting into Multiple Tables at Once	72
4.7 Blocking Inserts to Certain Columns	74
4.8 Modifying Records in a Table	75
4.9 Updating When Corresponding Rows Exist	77
4.10 Updating with Values from Another Table	78
4.11 Merging Records	81
4.12 Deleting All Records from a Table	83
4.13 Deleting Specific Records	83
4.14 Deleting a Single Record	84
4.15 Deleting Referential Integrity Violations	85
4.16 Deleting Duplicate Records	85
4.17 Deleting Records Referenced from Another Table	87
4.18 Summing Up	89
5. Metadata Queries.....	91
5.1 Listing Tables in a Schema	91
5.2 Listing a Table's Columns	93
5.3 Listing Indexed Columns for a Table	94
5.4 Listing Constraints on a Table	95
5.5 Listing Foreign Keys Without Corresponding Indexes	97

5.6 Using SQL to Generate SQL	100
5.7 Describing the Data Dictionary Views in an Oracle Database	102
5.8 Summing Up	103
6. Working with Strings.....	105
6.1 Walking a String	106
6.2 Embedding Quotes Within String Literals	108
6.3 Counting the Occurrences of a Character in a String	109
6.4 Removing Unwanted Characters from a String	110
6.5 Separating Numeric and Character Data	112
6.6 Determining Whether a String Is Alphanumeric	116
6.7 Extracting Initials from a Name	120
6.8 Ordering by Parts of a String	125
6.9 Ordering by a Number in a String	126
6.10 Creating a Delimited List from Table Rows	132
6.11 Converting Delimited Data into a Multivalued IN-List	136
6.12 Alphabetizing a String	141
6.13 Identifying Strings That Can Be Treated as Numbers	147
6.14 Extracting the nth Delimited Substring	153
6.15 Parsing an IP Address	160
6.16 Comparing Strings by Sound	162
6.17 Finding Text Not Matching a Pattern	164
6.18 Summing Up	167
7. Working with Numbers.....	169
7.1 Computing an Average	169
7.2 Finding the Min/Max Value in a Column	171
7.3 Summing the Values in a Column	173
7.4 Counting Rows in a Table	175
7.5 Counting Values in a Column	177
7.6 Generating a Running Total	178
7.7 Generating a Running Product	179
7.8 Smoothing a Series of Values	181
7.9 Calculating a Mode	182
7.10 Calculating a Median	185
7.11 Determining the Percentage of a Total	187
7.12 Aggregating Nullable Columns	190
7.13 Computing Averages Without High and Low Values	191
7.14 Converting Alphanumeric Strings into Numbers	193
7.15 Changing Values in a Running Total	196
7.16 Finding Outliers Using the Median Absolute Deviation	197
7.17 Finding Anomalies Using Benford's Law	201

7.18 Summing Up	203
8. Date Arithmetic	205
8.1 Adding and Subtracting Days, Months, and Years	205
8.2 Determining the Number of Days Between Two Dates	208
8.3 Determining the Number of Business Days Between Two Dates	210
8.4 Determining the Number of Months or Years Between Two Dates	215
8.5 Determining the Number of Seconds, Minutes, or Hours Between Two Dates	218
8.6 Counting the Occurrences of Weekdays in a Year	220
8.7 Determining the Date Difference Between the Current Record and the Next Record	231
8.8 Summing Up	237
9. Date Manipulation	239
9.1 Determining Whether a Year Is a Leap Year	240
9.2 Determining the Number of Days in a Year	246
9.3 Extracting Units of Time from a Date	249
9.4 Determining the First and Last Days of a Month	252
9.5 Determining All Dates for a Particular Weekday Throughout a Year	255
9.6 Determining the Date of the First and Last Occurrences of a Specific Weekday in a Month	261
9.7 Creating a Calendar	268
9.8 Listing Quarter Start and End Dates for the Year	281
9.9 Determining Quarter Start and End Dates for a Given Quarter	286
9.10 Filling in Missing Dates	293
9.11 Searching on Specific Units of Time	301
9.12 Comparing Records Using Specific Parts of a Date	302
9.13 Identifying Overlapping Date Ranges	305
9.14 Summing Up	311
10. Working with Ranges	313
10.1 Locating a Range of Consecutive Values	313
10.2 Finding Differences Between Rows in the Same Group or Partition	317
10.3 Locating the Beginning and End of a Range of Consecutive Values	323
10.4 Filling in Missing Values in a Range of Values	326
10.5 Generating Consecutive Numeric Values	330
10.6 Summing Up	333
11. Advanced Searching	335
11.1 Paginating Through a Result Set	335
11.2 Skipping n Rows from a Table	338

11.3 Incorporating OR Logic When Using Outer Joins	339
11.4 Determining Which Rows Are Reciprocals	341
11.5 Selecting the Top n Records	343
11.6 Finding Records with the Highest and Lowest Values	344
11.7 Investigating Future Rows	345
11.8 Shifting Row Values	347
11.9 Ranking Results	350
11.10 Suppressing Duplicates	351
11.11 Finding Knight Values	353
11.12 Generating Simple Forecasts	359
11.13 Summing Up	367
12. Reporting and Reshaping.....	369
12.1 Pivoting a Result Set into One Row	369
12.2 Pivoting a Result Set into Multiple Rows	372
12.3 Reverse Pivoting a Result Set	377
12.4 Reverse Pivoting a Result Set into One Column	379
12.5 Suppressing Repeating Values from a Result Set	382
12.6 Pivoting a Result Set to Facilitate Inter-Row Calculations	384
12.7 Creating Buckets of Data, of a Fixed Size	386
12.8 Creating a Predefined Number of Buckets	388
12.9 Creating Horizontal Histograms	390
12.10 Creating Vertical Histograms	392
12.11 Returning Non-GROUP BY Columns	394
12.12 Calculating Simple Subtotals	397
12.13 Calculating Subtotals for All Possible Expression Combinations	400
12.14 Identifying Rows That Are Not Subtotals	410
12.15 Using Case Expressions to Flag Rows	412
12.16 Creating a Sparse Matrix	414
12.17 Grouping Rows by Units of Time	416
12.18 Performing Aggregations over Different Groups/Partitions Simultaneously	420
12.19 Performing Aggregations over a Moving Range of Values	422
12.20 Pivoting a Result Set with Subtotals	429
12.21 Summing Up	434
13. Hierarchical Queries.....	435
13.1 Expressing a Parent-Child Relationship	436
13.2 Expressing a Child-Parent-Grandparent Relationship	440
13.3 Creating a Hierarchical View of a Table	444
13.4 Finding All Child Rows for a Given Parent Row	449
13.5 Determining Which Rows Are Leaf, Branch, or Root Nodes	450

13.6 Summing Up	458
14. Odds 'n' Ends.....	459
14.1 Creating Cross-Tab Reports Using SQL Server's PIVOT Operator	459
14.2 Unpivoting a Cross-Tab Report Using SQL Server's UNPIVOT Operator	461
14.3 Transposing a Result Set Using Oracle's MODEL Clause	463
14.4 Extracting Elements of a String from Unfixed Locations	467
14.5 Finding the Number of Days in a Year (an Alternate Solution for Oracle)	470
14.6 Searching for Mixed Alphanumeric Strings	472
14.7 Converting Whole Numbers to Binary Using Oracle	474
14.8 Pivoting a Ranked Result Set	477
14.9 Adding a Column Header into a Double Pivoted Result Set	481
14.10 Converting a Scalar Subquery to a Composite Subquery in Oracle	493
14.11 Parsing Serialized Data into Rows	495
14.12 Calculating Percent Relative to Total	500
14.13 Testing for Existence of a Value Within a Group	502
14.14 Summing Up	505
A. Window Function Refresher.....	507
B. Common Table Expressions.....	535
Index.....	539