

## Effective Sql 61 Specific Ways To Write Better Sql Effective Software Development

An industry consultant shares his most useful tips and tricks for advanced SQL programming to help the working programmer gain performance and work around system deficiencies.

Sams Teach Yourself SQL in 10 Minutes offers straightforward, practical answers when you need fast results. By working through the book's 22 lessons of 10 minutes or less, you'll learn what you need to know to take advantage of the SQL language. Lessons cover IBM DB2, Microsoft SQL Server and SQL Server Express, MariaDB, MySQL, Oracle and Oracle express, PostgreSQL, and SQLite. Full-color code examples help you understand how SQL statements are structured Tips point out shortcuts and solutions Cautions help you avoid common pitfalls Notes explain additional concepts, and provide additional information 10 minutes is all you need to learn how to... Use the major SQL statements Construct complex SQL statements using multiple clauses and operators Retrieve, sort, and format database contents Pinpoint the data you need using a variety of filtering techniques Use aggregate functions to summarize data Join two or more related tables Insert, update, and delete data Create and alter database tables Work with views, stored procedures, and more

A guide to SQL covers such topics as retrieving records, metadata queries, working with strings, data arithmetic, date manipulation, reporting and warehousing, and hierarchical queries.

Data science libraries, frameworks, modules, and toolkits are great for doing data science, but they're also a good way to dive into the discipline without actually understanding data science. In this book, you'll learn how many of the most fundamental data science tools and algorithms work by implementing them from scratch. If you have an aptitude for mathematics and some programming skills, author Joel Grus will help you get comfortable with the math and statistics at the core of data science, and with hacking skills you need to get started as a data scientist. Today's messy glut of data holds answers to questions no one's even thought to ask. This book provides you with the know-how to dig those answers out. Get a crash course in Python Learn the basics of linear algebra, statistics, and probability—and understand how and when they're used in data science Collect, explore, clean, munge, and manipulate data Dive into the fundamentals of machine learning Implement models such as k-nearest Neighbors, Naive Bayes, linear and logistic regression, decision trees, neural networks, and clustering Explore recommender systems, natural language processing, network analysis, MapReduce, and databases

The #1 Easy, Common-Sense Guide to SQL Queries—Updated with More Advanced Techniques and Solutions Foreword by Keith W. Hare, Vice Chair, USA SQL Standards Committee SQL Queries for Mere Mortals has earned worldwide praise as the clearest, simplest tutorial on writing effective queries with the latest SQL standards and database applications. Now, author John L. Viescas has updated this hands-on classic with even more advanced and valuable techniques. Step by step, Viescas guides you through creating reliable queries for virtually any current SQL-based database. He demystifies all aspects of SQL query writing, from simple data selection and filtering to joining multiple tables and modifying sets of data. Building on the basics, Viescas shows how to solve challenging real-world problems, including applying multiple complex conditions on one table, performing sophisticated logical evaluations, and using unlinked tables to think “outside the box.” In two brand-new chapters, you learn how to perform complex calculations on groups for sophisticated reporting, and how to partition data into windows for more flexible aggregation. Practice all you want with downloadable sample databases for today's versions of Microsoft Office Access, Microsoft SQL Server, and the open source MySQL and PostgreSQL databases. Whether you're a DBA, developer, user, or student, there's no better way to master SQL. Coverage includes: Getting started: understanding what relational databases are, and ensuring that your database structures are sound SQL basics: using SELECT statements, creating expressions, sorting information with ORDER BY, and filtering data using WHERE Summarizing and grouping data with GROUP BY and HAVING clauses Drawing data from multiple tables: using INNER JOIN, OUTER JOIN, and UNION operators, and working with subqueries Modifying data sets with UPDATE, INSERT, and DELETE statements Advanced queries: complex NOT and AND, conditions, if-then-else using CASE, unlinked tables, driver tables, and more NEW! Using advanced GROUP BY keywords to create subtotals, roll-ups, and more NEW! Applying window functions to answer more sophisticated questions, and gain deeper insight into your data Software-Independent Approach! If you work with database software such as Access, MS SQL Server, Oracle, DB2, MySQL, Ingres, or any other SQL-based program, this book could save you hours of time and aggravation—before you write a single query!

The need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational “NoSQL” databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage's pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Polyglot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access.

A poorly performing database application not only costs users time, but also has an impact on other applications running on the same computer or the same network. SQL Tuning provides an essential next step for SQL developers and database administrators who want to extend their SQL tuning expertise and get the most from their database applications. There are two basic issues to focus on when tuning SQL: how to find and interpret the execution plan of an SQL statement and how to change SQL to get a specific alternate execution plan. SQL Tuning provides answers to these questions and addresses a third issue that's even more important: how to find the optimal execution plan for the query to use. Author Dan Tow outlines a timesaving method he's developed for finding the optimum execution plan--rapidly and systematically--regardless of the complexity of the SQL or the database platform being used. You'll learn how to understand and control SQL execution plans and how to diagram SQL queries to deduce the best execution plan for a query. Key chapters in the book include exercises to reinforce the concepts you've learned. SQL Tuning concludes by addressing special concerns and unique solutions to "unsolvable problems." Whether you are a

programmer who develops SQL-based applications or a database administrator or other who troubleshoots poorly tuned applications, SQL Tuning will arm you with a reliable and deterministic method for tuning your SQL queries to gain optimal performance.

CD-ROM contains cross-referenced code.

Effective SQL61 Specific Ways to Write Better SQL Addison-Wesley Professional

Describes the basics of SQL, database design, and how to create a database using the SQL language.

Perform fast interactive analytics against different data sources using the Trino high-performance distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Trino. Initially developed by Facebook, open source Trino is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber, and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Trino query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Trino's use cases and learn about tools that will help you connect to Trino and query data Go deeper: Learn Trino's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Trino in production: Secure Trino, monitor workloads, tune queries, and connect more applications; learn how other organizations apply Trino

Prepare for Microsoft Exam 70-761—and help demonstrate your real-world mastery of SQL Server 2016 Transact-SQL data management, queries, and database programming. Designed for experienced IT professionals ready to advance their status, Exam Ref focuses on the critical-thinking and decision-making acumen needed for success at the MCSA level. Focus on the expertise measured by these objectives: • Filter, sort, join, aggregate, and modify data • Use subqueries, table expressions, grouping sets, and pivoting • Query temporal and non-relational data, and output XML or JSON • Create views, user-defined functions, and stored procedures • Implement error handling, transactions, data types, and nulls This Microsoft Exam Ref: • Organizes its coverage by exam objectives • Features strategic, what-if scenarios to challenge you • Assumes you have experience working with SQL Server as a database administrator, system engineer, or developer • Includes downloadable sample database and code for SQL Server 2016 SP1 (or later) and Azure SQL Database Querying Data with Transact-SQL About the Exam Exam 70-761 focuses on the skills and knowledge necessary to manage and query data and to program databases with Transact-SQL in SQL Server 2016. About Microsoft Certification Passing this exam earns you credit toward a Microsoft Certified Solutions Associate (MCSA) certification that demonstrates your mastery of essential skills for building and implementing on-premises and cloud-based databases across organizations. Exam 70-762 (Developing SQL Databases) is also required for MCSA: SQL 2016 Database Development certification. See full details at: [microsoft.com/learning](https://microsoft.com/learning)

Data Mining: Concepts and Techniques provides the concepts and techniques in processing gathered data or information, which will be used in various applications. Specifically, it explains data mining and the tools used in discovering knowledge from the collected data. This book is referred as the knowledge discovery from data (KDD). It focuses on the feasibility, usefulness, effectiveness, and scalability of techniques of large data sets. After describing data mining, this edition explains the methods of knowing, preprocessing, processing, and warehousing data. It then presents information about data warehouses, online analytical processing (OLAP), and data cube technology. Then, the methods involved in mining frequent patterns, associations, and correlations for large data sets are described. The book details the methods for data classification and introduces the concepts and methods for data clustering. The remaining chapters discuss the outlier detection and the trends, applications, and research frontiers in data mining. This book is intended for Computer Science students, application developers, business professionals, and researchers who seek information on data mining. Presents dozens of algorithms and implementation examples, all in pseudo-code and suitable for use in real-world, large-scale data mining projects Addresses advanced topics such as mining object-relational databases, spatial databases, multimedia databases, time-series databases, text databases, the World Wide Web, and applications in several fields Provides a comprehensive, practical look at the concepts and techniques you need to get the most out of your data

"Master every business SQL skill you need! Grouping, totaling, summaries, modifying databases, integrating data from multiple tables, and much more! Includes video introduction, 2+ hours of expert audio commentary, 200+ animated figures, 100+ self review questions, 100+ exercises, searching, hyperlinking, and more."--Container.

This book is an introduction and deep-dive into the many uses of dynamic SQL in Microsoft SQL Server. Dynamic SQL is key to large-scale searching based upon user-entered criteria. It's also useful in generating value-lists, in dynamic pivoting of data for business intelligence reporting, and for customizing database objects and querying their structure. Executing dynamic SQL is at the heart of applications such as business intelligence dashboards that need to be fluid and respond instantly to changing user needs as those users explore their data and view the results. Yet dynamic SQL is feared by many due to concerns over SQL injection attacks. Reading Dynamic SQL: Applications, Performance, and Security is your opportunity to learn and master an often misunderstood feature, including security and SQL injection. All aspects of security relevant to dynamic SQL are discussed in this book. You will learn many ways to save time and develop code more efficiently, and you will practice directly with security scenarios that threaten companies around the world every day. Dynamic SQL: Applications, Performance, and Security helps you bring the productivity and user-satisfaction of flexible and responsive applications to your organization safely and securely. Your organization's increased ability to respond to rapidly changing business scenarios will build competitive advantage in an increasingly crowded and competitive global marketplace. Discusses many applications of dynamic SQL, both simple and complex. Explains each example with demos that can be run at home and on your laptop. Helps you to identify when dynamic SQL

can offer superior performance. Pays attention to security and best practices to ensure safety of your data. What You Will Learn Build flexible applications that respond fast to changing business needs. Take advantage of unconventional but productive uses of dynamic SQL. Protect your data from attack through best-practices in your implementations. Know about SQL Injection and be confident in your defenses against it Run at high performance by optimizing dynamic SQL in your applications. Troubleshoot and debug dynamic SQL to ensure correct results. Who This Book is For Dynamic SQL: Applications, Performance, and Security is for developers and database administrators looking to hone and build their T-SQL coding skills. The book is ideal for advanced users wanting to plumb the depths of application flexibility and troubleshoot performance issues involving dynamic SQL. The book is also ideal for beginners wanting to learn what dynamic SQL is about and how it can help them deliver competitive advantage to their organizations.

SQL is a widely used to access most databases, therefore database developers and system administrators should be familiar with it. This hands-on SQL book will help beginner and intermediate users to write queries that apply complex conditions on a table. The book's unique side by side approach makes it easy for the reader to learn three major query languages in the IT industry. The author has over 20 years of experience in database design. KEY FEATURES: Contains numerous practical screenshots of Oracle SQL, T-SQL, MySQL statements and results. Shows the differences between Oracle SQL, T-SQL and MySQL side by side. Gives a real world experience for SQL developers and database administrators. Sample data is available to work on (available on our website).

Presents a guide to writing effective SQL queries, from simple data selection and filtering to joining multiple tables and modifying sets of data, with information on how to solve a variety of challenging SQL problems.

This pocket guide presents the most crucial information about SQL in a compact and easily accessible format, covering the four commonly used SQL variants--Oracle, IBM DB2, Microsoft SQL Server, and MySQL. Topics include: Data manipulation statements (SELECT, DELETE, INSERT, UPDATE, MERGE) and transaction control statements (START TRANSACTION, SAVEPOINT, COMMIT, ROLLBACK). Common SQL functions (date, numeric, math, trigonometric, string, conversion, aggregate) Such topics as literals, NULLs, CASE expressions, datatype conversion, regular expressions, grouping and summarizing data, joining tables, and writing queries (hierarchical, recursive, union, flashback) and subqueries. Instead of presenting complex and confusing syntax diagrams, the book teaches by example, showing the SQL statements and options that readers are most like to use. All example data is available on the O'Reilly web site. "If you need fast, accurate SQL information, with examples for multiple database engines, be sure to check out this book."--Chris Kempster, Senior DBA and author of SQL Server 2000 for the Oracle DBA, [www.chriskempster.com](http://www.chriskempster.com)

A complete guide to SQL\*Loader, a utility used to move data from external files into an Oracle database, offers step-by-step instruction in the various applications of SQL\*Loader, providing a task-oriented approach that covers the latest Oracle 8 and Oracle 8i features. Original. (Beginner/Intermediate)

Updated for the latest database management systems -- including MySQL 6.0, Oracle 11g, and Microsoft's SQL Server 2008 -- this introductory guide will get you up and running with SQL quickly. Whether you need to write database applications, perform administrative tasks, or generate reports, Learning SQL, Second Edition, will help you easily master all the SQL fundamentals. Each chapter presents a self-contained lesson on a key SQL concept or technique, with numerous illustrations and annotated examples. Exercises at the end of each chapter let you practice the skills you learn. With this book, you will: Move quickly through SQL basics and learn several advanced features Use SQL data statements to generate, manipulate, and retrieve data Create database objects, such as tables, indexes, and constraints, using SQL schema statements Learn how data sets interact with queries, and understand the importance of subqueries Convert and manipulate data with SQL's built-in functions, and use conditional logic in data statements Knowledge of SQL is a must for interacting with data. With Learning SQL, you'll quickly learn how to put the power and flexibility of this language to work.

Class-tested and coherent, this textbook teaches classical and web information retrieval, including web search and the related areas of text classification and text clustering from basic concepts. It gives an up-to-date treatment of all aspects of the design and implementation of systems for gathering, indexing, and searching documents; methods for evaluating systems; and an introduction to the use of machine learning methods on text collections. All the important ideas are explained using examples and figures, making it perfect for introductory courses in information retrieval for advanced undergraduates and graduate students in computer science. Based on feedback from extensive classroom experience, the book has been carefully structured in order to make teaching more natural and effective. Slides and additional exercises (with solutions for lecturers) are also available through the book's supporting website to help course instructors prepare their lectures.

Offers tips for improving the performance of any SQL database, no matter what the platform. Written for experienced database administrators familiar with SQL, the book identifies the similarities and differences of eight DBMSs, including Oracle 9i, IBM DB2 7.2, and Microsoft SQL server 2000. It provides strategies for refining sorts, subqueries, columns, tables, indexes, constraints, and locks. Annotation copyrighted by Book News, Inc., Portland, OR

Useful business analysis requires you to effectively transform data into actionable information. This book helps you use SQL and Excel to extract business information from relational databases and use that data to define business dimensions, store transactions about customers, produce results, and more. Each chapter explains when and why to perform a particular type of business analysis in order to obtain useful results, how to design and perform the analysis using SQL and Excel, and what the results should look like.

Languages may come and go, but the relational database endures. Learn how to use Ecto, the premier database library for Elixir, to connect your Elixir and Phoenix apps to databases. Get a firm handle on Ecto fundamentals with a module-by-module tour of the critical parts of Ecto. Then move on to more advanced topics and advice on best practices with a

series of recipes that provide clear, step-by-step instructions on scenarios commonly encountered by app developers. Co-authored by the creator of Ecto, this title provides all the essentials you need to use Ecto effectively. Elixir and Phoenix are taking the application development world by storm, and Ecto, the database library that ships with Phoenix, is going right along with them. There are plenty of examples that show you the basics, but to use Ecto to its full potential, you need to learn the library from the ground up. This definitive guide starts with a tour of the core features of Ecto - repos, queries, schemas, changesets, transactions - gradually building your knowledge with tasks of ever-increasing complexity. Along the way, you'll be learning by doing - a sample application handles all the boilerplate so you can focus on getting Ecto into your fingers. Build on that core knowledge with a series of recipes featuring more advanced topics. Change your pooling strategy to maximize your database's efficiency. Use nested associations to handle complex table relationships. Add streams to handle large result sets with ease. Based on questions from Ecto users, these recipes cover the most common situations developers run into. Whether you're new to Ecto, or already have an app in production, this title will give you a deeper understanding of how Ecto works, and help make your database code cleaner and more efficient. What You Need: To follow along with the book, you should have Erlang/OTP 19+ and Elixir 1.4+ installed. The book will guide you through setting up a sample application that integrates Ecto.

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

Effective SQL brings together the hands-on solutions and practical insights you need to solve a wide range of complex problems with SQL, and to design databases that make it far easier to manage data in the future. Leveraging the proven format of the best-selling Effective series, it focuses on providing clear, practical explanations, expert tips, and plenty of realistic examples -- all in full color. Drawing on their immense experience as consultants and instructors, three world-class database experts identify specific challenges, and distill each solution into five pages or less. Throughout, they provide well-annotated SQL code designed for all leading platforms, as well as code for specific implementations ranging from SQL Server to Oracle and MySQL, wherever these vary or permit you to achieve your goal more efficiently. Going beyond mere syntax, the authors also show how to avoid poor database design that makes it difficult to write effective SQL, how to improve suboptimal designs, and how to work around designs you can't change. You'll also find detailed sections on filtering and finding data, aggregation, subqueries, and metadata, as well as specific solutions for everything from listing products to scheduling events and defining data hierarchies. Simply put, if you already know the basics of SQL, Effective SQL will help you become a world-class SQL problem-solver.

The Essential Guide to SQL Statement Optimization Written by Oracle ACE Director and MySQL expert Ronald Bradford, Effective MySQL: Optimizing SQL Statements is filled with detailed explanations and practical examples that can be applied immediately to improve database and application performances. Featuring a step-by-step approach to SQL optimization, this Oracle Press book helps you to analyze and tune problematic SQL statements. Identify the essential analysis commands for gathering and diagnosing issues Learn how different index theories are applied and represented in MySQL Plan and execute informed SQL optimizations Create MySQL indexes to improve query performance Master the MySQL query execution plan Identify key configuration variables that impact SQL execution and performance Apply the SQL optimization lifecycle to capture, identify, confirm, analyze, and optimize SQL statements and verify the results Improve index utilization with covering indexes and partial indexes Learn hidden performance tips for improving index efficiency and simplifying SQL statements

SQL: The Ultimate Beginners Guide - Learn SQL Today Learning the SQL language can be laborious and tedious, but if you have genuine interest in learning a new language and updating your skills, it could be relatively easy. In this book, all the basic information that you need to learn as a beginner are presented. All you have to do is to apply them. This book will serve as an essential guide for you, as a SQL beginner. In addition, the concepts of SQL are laid out in a simple, concise language and instructions to help you learn the steps properly. Specific examples and sample tables is showcased to help you practice most of the SQL queries.

Tackle the toughest set-based querying and query tuning problems—guided by an author team with in-depth, inside knowledge of T-SQL. Deepen your understanding of architecture and internals—and gain practical approaches and advanced techniques to optimize your code's performance. Discover how to: Move from procedural programming to the language of sets and logic Optimize query tuning with a top-down methodology Assess algorithmic complexity to predict performance Compare data-aggregation techniques, including new grouping sets Manage data modification—insert, delete, update, merge—for performance Write more efficient queries against partitioned tables Work with graphs, trees, hierarchies, and recursive queries Plus—Use pure-logic puzzles to sharpen your problem-solving skills

Effectively query and modify data using Transact-SQL Master T-SQL fundamentals and write robust code for Microsoft SQL Server and Azure SQL Database. Itzik Ben-Gan explains key T-SQL concepts and helps you apply your knowledge

with hands-on exercises. The book first introduces T-SQL's roots and underlying logic. Next, it walks you through core topics such as single-table queries, joins, subqueries, table expressions, and set operators. Then the book covers more-advanced data-query topics such as window functions, pivoting, and grouping sets. The book also explains how to modify data, work with temporal tables, and handle transactions, and provides an overview of programmable objects. Microsoft Data Platform MVP Itzik Ben-Gan shows you how to: Review core SQL concepts and its mathematical roots Create tables and enforce data integrity Perform effective single-table queries by using the SELECT statement Query multiple tables by using joins, subqueries, table expressions, and set operators Use advanced query techniques such as window functions, pivoting, and grouping sets Insert, update, delete, and merge data Use transactions in a concurrent environment Get started with programmable objects—from variables and batches to user-defined functions, stored procedures, triggers, and dynamic SQL

Combines language tutorials with application design advice to cover the PHP server-side scripting language and the MySQL database engine.

Every software developer and IT professional understands the crucial importance of effective debugging. Often, debugging consumes most of a developer's workday, and mastering the required techniques and skills can take a lifetime. In *Effective Debugging*, Diomidis Spinellis helps experienced programmers accelerate their journey to mastery, by systematically categorizing, explaining, and illustrating the most useful debugging methods, strategies, techniques, and tools. Drawing on more than thirty-five years of experience, Spinellis expands your arsenal of debugging techniques, helping you choose the best approaches for each challenge. He presents vendor-neutral, example-rich advice on general principles, high-level strategies, concrete techniques, high-efficiency tools, creative tricks, and the behavioral traits associated with effective debugging. Spinellis's 66 expert techniques address every facet of debugging and are illustrated with step-by-step instructions and actual code. He addresses the full spectrum of problems that can arise in modern software systems, especially problems caused by complex interactions among components and services running on hosts scattered around the planet. Whether you're debugging isolated runtime errors or catastrophic enterprise system failures, this guide will help you get the job done—more quickly, and with less pain. Key features include High-level strategies and methods for addressing diverse software failures Specific techniques to apply when programming, compiling, and running code Better ways to make the most of your debugger General-purpose skills and tools worth investing in Advanced ideas and techniques for escaping dead-ends and the maze of complexity Advice for making programs easier to debug Specialized approaches for debugging multithreaded, asynchronous, and embedded code Bug avoidance through improved software design, construction, and management

Improve the performance of relational databases with indexes designed for today's hardware Over the last few years, hardware and software have advanced beyond all recognition, so it's hardly surprising that relational database performance now receives much less attention. Unfortunately, the reality is that the improved hardware hasn't kept pace with the ever-increasing quantity of data processed today. Although disk packing densities have increased enormously, making storage costs extremely low and sequential read very fast, random reads are still painfully slow. Many of the old design recommendations are therefore no longer valid—the optimal point of indexing has come a long way. Consequently many of the old problems haven't actually gone away—they have simply changed their appearance. This book provides an easy but effective approach to the design of indexes and tables. Using lots of examples and case studies, the authors describe how the DB2, Oracle, and SQL Server optimizers determine how to access data, and how CPU and response times for the resulting access paths can be quickly estimated. This enables comparisons to be made of the various designs, and helps you choose available choices for the most appropriate design. This book is intended for anyone who wants to understand the issues of SQL performance or how to design tables and indexes effectively. With this title, readers with many years of experience of relational systems will be able to better grasp the implications that have been brought into play by the introduction of new hardware.

Presents an instructional guide to SQL which uses humor and simple images to cover such topics as the structure of relational databases, simple and complex queries, creating multiple tables, and protecting important table data.

Demonstrates important concepts and offers working Transact-SQL code, covering data filtering, DDL, DML, statistical functions, runs and sequences, transactions, stored procedures and triggers, and performance tuning.

For all the buzz about trendy IT techniques, data processing is still at the core of our systems, especially now that enterprises all over the world are confronted with exploding volumes of data. Database performance has become a major headache, and most IT departments believe that developers should provide simple SQL code to solve immediate problems and let DBAs tune any bad SQL later. In *The Art of SQL*, author and SQL expert Stephane Faroult argues that this safe approach only leads to disaster. His insightful book, named after *Art of War* by Sun Tzu, contends that writing quick inefficient code is sweeping the dirt under the rug. SQL code may run for 5 to 10 years, surviving several major releases of the database management system and on several generations of hardware. The code must be fast and sound from the start, and that requires a firm understanding of SQL and relational theory. *The Art of SQL* offers best practices that teach experienced SQL users to focus on strategy rather than specifics. Faroult's approach takes a page from Sun Tzu's classic treatise by viewing database design as a military campaign. You need knowledge, skills, and talent. Talent can't be taught, but every strategist from Sun Tzu to modern-day generals believed that it can be nurtured through the experience of others. They passed on their experience acquired in the field through basic principles that served as guiding stars amid the sound and fury of battle. This is what Faroult does with SQL. Like a successful battle plan, good architectural choices are based on contingencies. What if the volume of this or that table increases unexpectedly? What if, following a merger, the number of users doubles? What if you want to keep several years of data online? Faroult's way of looking at SQL performance may be unconventional and unique, but he's deadly serious about writing good SQL and using SQL well. *The Art of SQL* is not a cookbook, listing problems and giving recipes. The aim is to get you—and your manager—to raise good questions.

*Practical SQL* is an approachable and fast-paced guide to SQL (Structured Query Language), the standard programming language for defining, organizing, and exploring data in relational databases. The book focuses on using SQL to find the story your data tells, with the popular open-source database PostgreSQL and the pgAdmin interface as its primary tools. You'll first cover the fundamentals of databases and the SQL language, then build skills by analyzing data from the U.S. Census and other federal and state government agencies. With exercises and real-world examples in each chapter, this book will teach even those who have never programmed before all the tools necessary to build powerful databases and access information quickly and efficiently. You'll learn how to: - Create databases and related tables using your own data - Define the right data types for your information - Aggregate, sort, and filter data to find patterns - Use basic math and advanced statistical functions - Identify errors in data and clean them up - Import and export data using delimited text files - Write queries

for geographic information systems (GIS) - Create advanced queries and automate tasks Learning SQL doesn't have to be dry and complicated. Practical SQL delivers clear examples with an easy-to-follow approach to teach you the tools you need to build and manage your own databases. This book uses PostgreSQL, but the SQL syntax is applicable to many database applications, including Microsoft SQL Server and MySQL.

If you have mastered the fundamentals of the PL/SQL language and are now looking for an in-depth, practical guide to solving real problems with PL/SQL stored procedures, then this is the book for you.

Design and write simple and efficient T-SQL code in SQL Server 2019 and beyond. Writing T-SQL that pulls back correct results can be challenging. This book provides the help you need in writing T-SQL that performs fast and is easy to maintain. You also will learn how to implement version control, testing, and deployment strategies. Hands-on examples show modern T-SQL practices and provide straightforward explanations. Attention is given to selecting the right data types and objects when designing T-SQL solutions. Author Elizabeth Noble teaches you how to improve your T-SQL performance through good design practices that benefit programmers and ultimately the users of the applications. You will know the common pitfalls of writing T-SQL and how to avoid those pitfalls going forward. What You Will Learn Choose correct data types and database objects when designing T-SQL Write T-SQL that searches data efficiently and uses hardware effectively Implement source control and testing methods to streamline the deployment process Design T-SQL that can be enhanced or modified with less effort Plan for long-term data management and storage Who This Book Is For Database developers who want to improve the efficiency of their applications, and developers who want to solve complex query and data problems more easily by writing T-SQL that performs well, brings back correct results, and is easy for other developers to understand and maintain

[Copyright: 9a81eed8daa719595150cd108f0fb101](#)