
Read Book SQL Antipatterns Avoiding The Pitfalls Of Database Programming Pragmatic Programmers

Right here, we have countless book **SQL Antipatterns Avoiding The Pitfalls Of Database Programming Pragmatic Programmers** and collections to check out. We additionally offer variant types and afterward type of the books to browse. The all right book, fiction, history, novel, scientific research, as well as various further sorts of books are readily open here.

As this SQL Antipatterns Avoiding The Pitfalls Of Database Programming Pragmatic Programmers, it ends taking place living thing one of the favored books SQL Antipatterns Avoiding The Pitfalls Of Database Programming Pragmatic Programmers collections that we have. This is why you remain in the best website to see the amazing ebook to have.

657ZRO - PETERSEN NATHAN

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.

Are you an SQL programmer that, like many, came to SQL after learning and writing procedural or object-oriented code? Or have switched jobs to where a different brand of SQL is being used, or maybe even been told to learn SQL yourself? If even one answer is yes, then you need this book. A "Manual of Style" for the SQL programmer, this book is a collection of heuristics and rules, tips, and tricks that will help you improve SQL programming style and proficiency, and for formatting and writing portable, readable, maintainable SQL code. Based on many years of experience consulting in SQL shops, and gathering questions and resolving his

students' SQL style issues, Joe Celko can help you become an even better SQL programmer. Help you write Standard SQL without an accent or a dialect that is used in another programming language or a specific flavor of SQL, code that can be maintained and used by other people. Enable you to give your group a coding standard for internal use, to enable programmers to use a consistent style. Give you the mental tools to approach a new problem with SQL as your tool, rather than another programming language — one that someone else might not know!

Learn everything you need to know to build efficient SQL queries using this easy-to-follow beginner's guide Key FeaturesExplore all SQL statements in depth using a variety of examplesGet to grips with database querying, data aggregate, manipulation, and much moreUnderstand how to explore and process data of varying complexity to tell a storyBook Description SQL is a powerful querying

language that's used to store, manipulate, and retrieve data, and it is one of the most popular languages used by developers to query and analyze data efficiently. If you're looking for a comprehensive introduction to SQL, *Learn SQL Database Programming* will help you to get up to speed with using SQL to streamline your work in no time. Starting with an overview of relational database management systems, this book will show you how to set up and use MySQL Workbench and design a database using practical examples. You'll also discover how to query and manipulate data with SQL programming using MySQL Workbench. As you advance, you'll create a database, query single and multiple tables, and modify data using SQL querying. This SQL book covers advanced SQL techniques, including aggregate functions, flow control statements, error handling, and subqueries, and helps you process your data to present your findings. Finally, you'll implement best practices for writing SQL and designing indexes and tables. By the end of this SQL programming book, you'll have gained the confidence to use SQL queries to retrieve and manipulate data. What you will learn: Install, configure, and use MySQL Workbench to restore a database; Explore different data types such as string, numeric, and date and time; Query a single table using the basic SQL SELECT statement and the FROM, WHERE, and ORDER BY clauses; Query multiple tables by understanding various types of table relationships; Modify data in tables using the INSERT, UPDATE, and DELETE statements; Use aggregate functions to group and summarize data; Detect bad data, duplicates, and irrelevant values while processing data; Who this book is for: This book is for business analysts, SQL developers, database administrators, and students learning SQL. If you want to learn how to query and ma-

nipulate SQL data for database administration tasks or simply extract and organize relevant data for analysis, you'll find this book useful. No prior SQL experience is required.

Explore the world of .NET design patterns and bring the benefits that the right patterns can offer to your toolkit today About This Book Dive into the powerful fundamentals of .NET framework for software development The code is explained piece by piece and the application of the pattern is also showcased. This fast-paced guide shows you how to implement the patterns into your existing applications Who This Book Is For This book is for those with familiarity with .NET development who would like to take their skills to the next level and be in the driver's seat when it comes to modern development techniques. Basic object-oriented C# programming experience and an elementary familiarity with the .NET framework library is required. What You Will Learn Put patterns and pattern catalogs into the right perspective Apply patterns for software development under C#/.NET Use GoF and other patterns in real-life development scenarios Be able to enrich your design vocabulary and well articulate your design thoughts Leverage object/functional programming by mixing OOP and FP Understand the reactive programming model using Rx and RxJs Writing compositional code using C# LINQ constructs Be able to implement concurrent/parallel programming techniques using idioms under .NET Avoiding pitfalls when creating compositional, readable, and maintainable code using imperative, functional, and reactive code. In Detail Knowing about design patterns enables developers to improve their code base, promoting code reuse and making their design more robust. This book focuses on the practical aspects of programming in .NET. You will learn about some of

the relevant design patterns (and their application) that are most widely used. We start with classic object-oriented programming (OOP) techniques, evaluate parallel programming and concurrency models, enhance implementations by mixing OOP and functional programming, and finally to the reactive programming model where functional programming and OOP are used in synergy to write better code. Throughout this book, we'll show you how to deal with architecture/design techniques, GoF patterns, relevant patterns from other catalogs, functional programming, and reactive programming techniques. After reading this book, you will be able to convincingly leverage these design patterns (factory pattern, builder pattern, prototype pattern, adapter pattern, facade pattern, decorator pattern, observer pattern and so on) for your programs. You will also be able to write fluid functional code in .NET that would leverage concurrency and parallelism! Style and approach This tutorial-based book takes a step-by-step approach. It covers the major patterns and explains them in a detailed manner along with code examples.

Develop microservice-based enterprise applications with expert guidance to avoid failures and technological debt with the help of real-world examples
Key Features
Implement the right microservices adoption strategy to transition from monoliths to microservices
Explore real-world use cases that explain anti-patterns and alternative practices in microservices development
Discover proven recommendations for avoiding architectural mistakes when designing microservices
Book Description
Microservices have been widely adopted for designing distributed enterprise apps that are flexible, robust, and fine-grained into services that

are independent of each other. There has been a paradigm shift where organizations are now either building new apps on microservices or transforming existing monolithic apps into microservices-based architecture. This book explores the importance of anti-patterns and the need to address flaws in them with alternative practices and patterns. You'll identify common mistakes caused by a lack of understanding when implementing microservices and cover topics such as organizational readiness to adopt microservices, domain-driven design, and resiliency and scalability of microservices. The book further demonstrates the anti-patterns involved in re-platforming brownfield apps and designing distributed data architecture. You'll also focus on how to avoid communication and deployment pitfalls and understand cross-cutting concerns such as logging, monitoring, and security. Finally, you'll explore testing pitfalls and establish a framework to address isolation, autonomy, and standardization. By the end of this book, you'll have understood critical mistakes to avoid while building microservices and the right practices to adopt early in the product life cycle to ensure the success of a microservices initiative. What you will learn
Discover the responsibilities of different individuals involved in a microservices initiative
Avoid the common mistakes in architecting microservices for scalability and resiliency
Understand the importance of domain-driven design when developing microservices
Identify the common pitfalls involved in migrating monolithic applications to microservices
Explore communication strategies, along with their potential drawbacks and alternatives
Discover the importance of adopting governance, security, and monitoring
Understand the role of CI/CD and testing
Who this book is for
This practical microservices book is for software archi-

itects, solution architects, and developers involved in designing microservices architecture and its development, who want to gain insights into avoiding pitfalls and drawbacks in distributed applications, and save time and money that might otherwise get wasted if microservices designs fail. Working knowledge of microservices is assumed to get the most out of this book.

The vast majority of software applications use relational databases that virtually every application developer must work with. This book introduces you to database design, whether you're a DBA or database developer. You'll discover what databases are, their goals, and why proper design is necessary to achieve those goals. Additionally, you'll master how to structure the database so it gives good performance while minimizing the chance for error. You will learn how to decide what should be in a database to meet the application's requirements.

Business rules management system (BRMS) is a software tools that work alongside enterprise IT applications. It enables enterprises to automate decision-making processes typically consisting of separate business rules authoring and rules execution applications. This proposed title brings together the following key ideas in modern enterprise system development best practice. The need for service-oriented architecture (SOA). How the former depends on component-based development (CBD). Database-centred approaches to business rules (inc. GUIDES). Knowledge-based approaches to business rules. Using patterns to design and develop business rules management systems Ian Graham is an industry consultant with over 20 years. He is recognized internationally as an authority on business modelling, object-oriented software development methods and expert systems.

He has a significant public presence, being associated with both UK and international professional organizations, and is frequently quoted in the IT and financial press.

What others in the trenches say about *The Pragmatic Programmer*... "The cool thing about this book is that it's great for keeping the programming process fresh. The book helps you to continue to grow and clearly comes from people who have been there."

—Kent Beck, author of *Extreme Programming Explained: Embrace Change*

"I found this book to be a great mix of solid advice and wonderful analogies!" —Martin Fowler, author of *Refactoring and UML Distilled*

"I would buy a copy, read it twice, then tell all my colleagues to run out and grab a copy. This is a book I would never loan because I would worry about it being lost."

—Kevin Ruland, Management Science, MSG-Logistics

"The wisdom and practical experience of the authors is obvious. The topics presented are relevant and useful.... By far its greatest strength for me has been the outstanding analogies—tracer bullets, broken windows, and the fabulous helicopter-based explanation of the need for orthogonality, especially in a crisis situation. I have little doubt that this book will eventually become an excellent source of useful information for journeymen programmers and expert mentors alike."

—John Lakos, author of *Large-Scale C++ Software Design*

"This is the sort of book I will buy a dozen copies of when it comes out so I can give it to my clients."

—Eric Vought, Software Engineer

"Most modern books on software development fail to cover the basics of what makes a great software developer, instead spending their time on syntax or technology where in reality the greatest leverage possible for any software team is in having talented developers who really know their craft well. An excel-

lent book.” —Pete McBreen, Independent Consultant “Since reading this book, I have implemented many of the practical suggestions and tips it contains. Across the board, they have saved my company time and money while helping me get my job done quicker! This should be a desktop reference for everyone who works with code for a living.” —Jared Richardson, Senior Software Developer, iRenaissance, Inc. “I would like to see this issued to every new employee at my company....” —Chris Cleeland, Senior Software Engineer, Object Computing, Inc. “If I’m putting together a project, it’s the authors of this book that I want. . . . And failing that I’d settle for people who’ve read their book.” —Ward Cunningham Straight from the programming trenches, *The Pragmatic Programmer* cuts through the increasing specialization and technicalities of modern software development to examine the core process—taking a requirement and producing working, maintainable code that delights its users. It covers topics ranging from personal responsibility and career development to architectural techniques for keeping your code flexible and easy to adapt and reuse. Read this book, and you'll learn how to Fight software rot; Avoid the trap of duplicating knowledge; Write flexible, dynamic, and adaptable code; Avoid programming by coincidence; Bullet-proof your code with contracts, assertions, and exceptions; Capture real requirements; Test ruthlessly and effectively; Delight your users; Build teams of pragmatic programmers; and Make your developments more precise with automation. Written as a series of self-contained sections and filled with entertaining anecdotes, thoughtful examples, and interesting analogies, *The Pragmatic Programmer* illustrates the best practices and major pitfalls of many different aspects of software development.

Whether you're a new coder, an experienced programmer, or a manager responsible for software projects, use these lessons daily, and you'll quickly see improvements in personal productivity, accuracy, and job satisfaction. You'll learn skills and develop habits and attitudes that form the foundation for long-term success in your career. You'll become a Pragmatic Programmer.

Provides information on the architecture of the T-SQL programming language to create scalable code.

T-SQL insiders help you tackle your toughest queries and query-tuning problems Squeeze maximum performance and efficiency from every T-SQL query you write or tune. Four leading experts take an in-depth look at T-SQL's internal architecture and offer advanced practical techniques for optimizing response time and resource usage. Emphasizing a correct understanding of the language and its foundations, the authors present unique solutions they have spent years developing and refining. All code and techniques are fully updated to reflect new T-SQL enhancements in Microsoft SQL Server 2014 and SQL Server 2012. Write faster, more efficient T-SQL code: Move from procedural programming to the language of sets and logic Master an efficient top-down tuning methodology Assess algorithmic complexity to predict performance Compare data aggregation techniques, including new grouping sets Efficiently perform data-analysis calculations Make the most of T-SQL's optimized bulk import tools Avoid date/time pitfalls that lead to buggy, poorly performing code Create optimized BI statistical queries without additional software Use programmable objects to accelerate queries Unlock major performance improvements with In-Memory OLTP Master useful and ele-

gant approaches to manipulating graphs About This Book For experienced T-SQL practitioners Includes coverage updated from Inside Microsoft SQL Server 2008 T-SQL Querying and Inside Microsoft SQL Server 2008 T-SQL Programming Valuable to developers, DBAs, BI professionals, and data scientists Covers many MCSE 70-464 and MCSA/MCSE 70-461 exam topics

The First Guide to Scrum-Based Agile Product Management In Agile Product Management with Scrum, leading Scrum consultant Roman Pichler uses real-world examples to demonstrate how product owners can create successful products with Scrum. He describes a broad range of agile product management practices, including making agile product discovery work, taking advantage of emergent requirements, creating the minimal marketable product, leveraging early customer feedback, and working closely with the development team. Benefitting from Pichler's extensive experience, you'll learn how Scrum product ownership differs from traditional product management and how to avoid and overcome the common challenges that Scrum product owners face. Coverage includes Understanding the product owner's role: what product owners do, how they do it, and the surprising implications Envisioning the product: creating a compelling product vision to galvanize and guide the team and stakeholders Grooming the product backlog: managing the product backlog effectively even for the most complex products Planning the release: bringing clarity to scheduling, budgeting, and functionality decisions Collaborating in sprint meetings: understanding the product owner's role in sprint meetings, including the dos and don'ts Transitioning into product ownership: succeeding as a product owner and establishing the role in the enterprise This book is an indis-

pensable resource for anyone who works as a product owner, or expects to do so, as well as executives and coaches interested in establishing agile product management.

When it comes to choosing, using, and maintaining a database, understanding its internals is essential. But with so many distributed databases and tools available today, it's often difficult to understand what each one offers and how they differ. With this practical guide, Alex Petrov guides developers through the concepts behind modern database and storage engine internals. Throughout the book, you'll explore relevant material gleaned from numerous books, papers, blog posts, and the source code of several open source databases. These resources are listed at the end of parts one and two. You'll discover that the most significant distinctions among many modern databases reside in subsystems that determine how storage is organized and how data is distributed. This book examines: Storage engines: Explore storage classification and taxonomy, and dive into B-Tree-based and immutable Log Structured storage engines, with differences and use-cases for each Storage building blocks: Learn how database files are organized to build efficient storage, using auxiliary data structures such as Page Cache, Buffer Pool and Write-Ahead Log Distributed systems: Learn step-by-step how nodes and processes connect and build complex communication patterns Database clusters: Which consistency models are commonly used by modern databases and how distributed storage systems achieve consistency

This book presents recent research work and results in the area of communication and information technologies. The book in-

cludes the main results of the 11th International Conference on Computing and Information Technology (IC2IT) held during July 2nd-3rd, 2015 in Bangkok, Thailand. The book is divided into the two main parts Data Mining and Machine Learning as well as Data Network and Communications. New algorithms and methods of data mining are discussed as well as innovative applications and state-of-the-art technologies on data mining, machine learning and data networking.

The Pragmatic Programmers classic is back! Freshly updated for modern software development, Pragmatic Unit Testing in Java 8 With JUnit teaches you how to write and run easily maintained unit tests in JUnit with confidence. You'll learn mnemonics to help you know what tests to write, how to remember all the boundary conditions, and what the qualities of a good test are. You'll see how unit tests can pay off by allowing you to keep your system code clean, and you'll learn how to handle the stuff that seems too tough to test. Pragmatic Unit Testing in Java 8 With JUnit steps you through all the important unit testing topics. If you've never written a unit test, you'll see screen shots from Eclipse, IntelliJ IDEA, and NetBeans that will help you get past the hard part-getting set up and started. Once past the basics, you'll learn why you want to write unit tests and how to effectively use JUnit. But the meaty part of the book is its collected unit testing wisdom from people who've been there, done that on production systems for at least 15 years: veteran author and developer Jeff Langr, building on the wisdom of Pragmatic Programmers Andy Hunt and Dave Thomas. You'll learn: How to craft your unit tests to minimize your effort in maintaining them. How to use unit tests to help keep your system clean. How to test the tough stuff. Mem-

orable mnemonics to help you remember what's important when writing unit tests. How to help your team reap and sustain the benefits of unit testing. You won't just learn about unit testing in theory--you'll work through numerous code examples. When it comes to programming, hands-on is the only way to learn!

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.

Six-Step Relational Database Design™ bridges the gaps between database theory, database modeling, and database implementation by outlining a simple but reliable six-step process for accurately modeling user data on a Crow's Foot Relational Model Diagram, and then demonstrating how to implement this model on any relational database management system. The second edition contains a new chapter on implementation that goes through the steps necessary to implement each of the case studies on a relational database management system, clearly relating the design to implementation and database theory. In addition, questions are also included at the end of each of the six steps and one of the previous case studies has been replaced, making the case study selection more diverse. Six-Step Relational Database Design™ uses three case studies and starts with a statement of the problem by the client and then goes through the six steps necessary to create a reliable and accurate data model of the client's business requirements. This model can then be used to implement the database on any relational database management system. Six-Step Relational Database Design™ should be used as a handbook for students and professionals in the software-development field. The technique described in this book can be used by

students for quickly developing relational databases for their applications, and by professionals for developing sturdy, reliable, and accurate relational database models for their software applications.

SQL is the ubiquitous language for software developers working with structured data. Most developers who rely on SQL are experts in their favorite language (such as Java, Python, or Go), but they're not experts in SQL. They often depend on antipatterns - solutions that look right but become increasingly painful to work with as you uncover their hidden costs. Learn to identify and avoid many of these common blunders. Refactor an inherited nightmare into a data model that really works. Updated for the current versions of MySQL and Python, this new edition adds a dozen brand new mini-antipatterns for quick wins. No matter which platform, framework, or language you use, the database is the foundation of your application, and the SQL database language is the standard for working with it. Antipatterns are solutions that look simple at the surface, but soon mire you down with needless work. Learn to identify these traps, and craft better solutions for the often-asked questions in this book. Avoid the mistakes that lead to poor performance and quality, and master the principles that make SQL a powerful and flexible tool for handling data and logic. Dive deep into SQL and database design, and learn to recognize the most common missteps made by software developers in database modeling, SQL query logic, and code design of data-driven applications. See practical examples of misconceptions about SQL that can lure software projects astray. Find the greatest value in each group of data. Understand why an

intersection table may be your new best friend. Store passwords securely and don't reinvent the wheel. Handle NULL values like a pro. Defend your web applications against the security weakness of SQL injection. Use SQL the right way - it can save you from headaches and needless work, and let your application really shine! What You Need: The SQL examples use the MySQL 8.0 flavor, but other popular brands of RDBMS are mentioned. Other code examples use Python 3.9+ or Ruby 2.7+.

API Design for C++ provides a comprehensive discussion of Application Programming Interface (API) development, from initial design through implementation, testing, documentation, release, versioning, maintenance, and deprecation. It is the only book that teaches the strategies of C++ API development, including interface design, versioning, scripting, and plug-in extensibility. Drawing from the author's experience on large scale, collaborative software projects, the text offers practical techniques of API design that produce robust code for the long term. It presents patterns and practices that provide real value to individual developers as well as organizations. API Design for C++ explores often overlooked issues, both technical and non-technical, contributing to successful design decisions that product high quality, robust, and long-lived APIs. It focuses on various API styles and patterns that will allow you to produce elegant and durable libraries. A discussion on testing strategies concentrates on automated API testing techniques rather than attempting to include end-user application testing techniques such as GUI testing, system testing, or manual testing. Each concept is illustrated with extensive C++ code examples, and fully functional examples and working source code for experimentation are available online. This book will be helpful

to new programmers who understand the fundamentals of C++ and who want to advance their design skills, as well as to senior engineers and software architects seeking to gain new expertise to complement their existing talents. Three specific groups of readers are targeted: practicing software engineers and architects, technical managers, and students and educators. The only book that teaches the strategies of C++ API development, including design, versioning, documentation, testing, scripting, and extensibility. Extensive code examples illustrate each concept, with fully functional examples and working source code for experimentation available online. Covers various API styles and patterns with a focus on practical and efficient designs for large-scale long-term projects.

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.

This indispensable SQL reference book is the first of its kind to leverage the benefits of design patterns to relational database SQL queries; all common SQL structures and design patterns are clearly categorized and described. Emphasizing the theoretical foundation for almost every type of SQL query problem, accompanying figures are included to help visualize the problem. Because SQL is a declarative language there are many ways to write any SQL query and professional database programmers must understand the correct way to write SQL for complicated database queries, and managers must institute formal SQL coding standards to improve productivity and maintainability. The SQL design patterns in this resource greatly improve the quality and productivity of systems development projects by forming a "best practices" foundation for all relational database queries.

Describes ways to incorporate domain modeling into software development.

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. Because databases often stay in production for decades, careful

design is critical to making the database serve the needs of your users over years, and to avoid subtle errors or performance problems. In this book, C.J. Date, a leading exponent of relational databases, lays out the principles of good database design.

Although MySQL's source code is open in the sense of being publicly available, it's essentially closed to you if you don't understand it. In this book, Sasha Pachev -- a former member of the MySQL Development Team -- provides a comprehensive tour of MySQL 5 that shows you how to figure out the inner workings of this powerful database. You'll go right to heart of the database to learn how data structures and convenience functions operate, how to add new storage engines and configuration options, and much more. The core of *Understanding MySQL Internals* begins with an Architecture Overview that provides a brief introduction of how the different components of MySQL work together. You then learn the steps for setting up a working compilable copy of the code that you can change and test at your pleasure. Other sections of the book cover: Core server classes, structures, and API The communication protocol between the client and the server Configuration variables, the controls of the server; includes a tutorial on how to add your own Thread-based request handling -- understanding threads and how they are used in MySQL An overview of MySQL storage engines The storage engine interface for integrating third-party storage engines The table lock manager The parser and optimizer for improving MySQL's performance Integrating a transactional storage engine into MySQL The internals of replication *Understanding MySQL Internals* provides unprecedented opportunities for developers, DBAs, database application programmers, IT departments, software vendors, and computer science

students to learn about the inner workings of this enterprise-proven database. With this book, you will soon reach a new level of comprehension regarding database development that will enable you to accomplish your goals. It's your guide to discovering and improving a great database.

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.

Build a core level of competency in SQL so you can recognize the

parts of queries and write simple SQL statements. SQL knowledge is essential for anyone involved in programming, data science, and data management. This book covers features of SQL that are standardized and common across most database vendors. You will gain a base of knowledge that will prepare you to go deeper into the specifics of any database product you might encounter. Examples in the book are worked in PostgreSQL and SQLite, but the bulk of the examples are platform agnostic and will work on any database platform supporting SQL. Early in the book you learn about table design, the importance of keys as row identifiers, and essential query operations. You then move into more advanced topics such as grouping and summarizing, creating calculated fields, joining data from multiple tables when it makes business sense to do so, and more. Throughout the book, you are exposed to a set-based approach to the language and are provided a good grounding in subtle but important topics such as the effects of null value on query results. With the explosion of data science, SQL has regained its prominence as a top skill to have for technologists and decision makers worldwide. SQL Primer will guide you from the very basics of SQL through to the mainstream features you need to have a solid, working knowledge of this important, data-oriented language. What You'll Learn Create and populate your own database tables Read SQL queries and understand what they are doing Execute queries that get correct results Bring together related rows from multiple tables Group and sort data in support of reporting applications Get a grip on nulls, normalization, and other key concepts Employ subqueries, unions, and other advanced features Who This Book Is For Anyone new to SQL who is looking for step-by-step guidance toward understand-

ing and writing SQL queries. The book is aimed at those who encounter SQL statements often in their work, and provides a sound baseline useful across all SQL database systems. Programmers, database managers, data scientists, and business analysts all can benefit from the baseline of SQL knowledge provided in this book. Summary Go in Practice guides you through 70 real-world techniques in key areas like package management, microservice communication, and more. Following a cookbook-style Problem/Solution/Discussion format, this practical handbook builds on the foundational concepts of the Go language and introduces specific strategies you can use in your day-to-day applications. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Go may be the perfect systems language. Built with simplicity, concurrency, and modern applications in mind, Go provides the core tool set for rapidly building web, cloud, and systems applications. If you know a language like Java or C#, it's easy to get started with Go; the trick is finding the practical dirt-under-the-fingernails techniques that you need to build production-ready code. About the Book Go in Practice guides you through dozens of real-world techniques in key areas. Following a cookbook-style Problem/Solution/Discussion format, this practical handbook builds on the foundational concepts of the Go language and introduces specific strategies you can use in your day-to-day applications. You'll learn techniques for building web services, using Go in the cloud, testing and debugging, routing, network applications, and much more. After finishing this book, you will be ready to build sophisticated cloud-native Go applications. What's Inside Dozens of specific, practical Golang techniques Using Go for devops and clou-

ops Writing RESTful web services and microservices Practical web dev techniques About the Reader Written for experienced developers who have already started exploring Go and want to use it effectively in a production setting. About the Authors Matt Farina is a software architect at Deis. Matt Butcher is a Principal Engineer in the Advanced Technology Group at Hewlett Packard Enterprise. They are both authors, speakers, and regular open source contributors. Table of Contents PART 1 - BACKGROUND AND FUNDAMENTALS Getting into Go A solid foundation Concurrency in Go PART 2 - WELL-ROUNDED APPLICATIONS Handling errors and panic Debugging and testing PART 3 - AN INTERFACE FOR YOUR APPLICATIONS HTML and email template patterns Serving and receiving assets and forms Working with web services PART 4 - TAKING YOUR APPLICATIONS TO THE CLOUD Using the cloud Communication between cloud services Reflection and code generation Mathematics is beautiful--and it can be fun and exciting as well as practical. Good Math is your guide to some of the most intriguing topics from two thousand years of mathematics: from Egyptian fractions to Turing machines; from the real meaning of numbers to proof trees, group symmetry, and mechanical computation. If you've ever wondered what lay beyond the proofs you struggled to complete in high school geometry, or what limits the capabilities of computer on your desk, this is the book for you. Why do Roman numerals persist? How do we know that some infinities are larger than others? And how can we know for certain a program will ever finish? In this fast-paced tour of modern and not-so-modern math, computer scientist Mark Chu-Carroll explores some of the greatest breakthroughs and disappointments of more than two thousand years of mathematical thought. There is joy and

beauty in mathematics, and in more than two dozen essays drawn from his popular "Good Math" blog, you'll find concepts, proofs, and examples that are often surprising, counterintuitive, or just plain weird. Mark begins his journey with the basics of numbers, with an entertaining trip through the integers and the natural, rational, irrational, and transcendental numbers. The voyage continues with a look at some of the oddest numbers in mathematics, including zero, the golden ratio, imaginary numbers, Roman numerals, and Egyptian and continuing fractions. After a deep dive into modern logic, including an introduction to linear logic and the logic-savvy Prolog language, the trip concludes with a tour of modern set theory and the advances and paradoxes of modern mechanical computing. If your high school or college math courses left you grasping for the inner meaning behind the numbers, Mark's book will both entertain and enlighten you.

Illustrating some of the most common misconceptions and pitfalls software developers face using relational databases, this book helps readers use a database to produce the most efficient results, and turn sluggish, inflexible code into high-quality, reliable solutions.

Troubleshoot query performance issues, identify anti-patterns in code, and write efficient T-SQL queries

Key Features

- Discover T-SQL functionalities and services that help you interact with relational databases
- Understand the roles, tasks and responsibilities of a T-SQL developer
- Explore solutions for carrying out database querying tasks, database administration, and troubleshooting

Book Description

Transact-SQL (T-SQL) is Microsoft's proprietary extension to the SQL language that is used with Microsoft SQL

Server and Azure SQL Database. This book will be a useful guide to learning the art of writing efficient T-SQL code in modern SQL Server versions, as well as the Azure SQL Database. The book will get you started with query processing fundamentals to help you write powerful, performant T-SQL queries. You will then focus on query execution plans and learn how to leverage them for troubleshooting. In the later chapters, you will learn how to identify various T-SQL patterns and anti-patterns. This will help you analyze execution plans to gain insights into current performance, and determine whether or not a query is scalable. You will also learn to build diagnostic queries using dynamic management views (DMVs) and dynamic management functions (DMFs) to address various challenges in T-SQL execution. Next, you will study how to leverage the built-in tools of SQL Server to shorten the time taken to address query performance and scalability issues. In the concluding chapters, the book will guide you through implementing various features, such as Extended Events, Query Store, and Query Tuning Assistant using hands-on examples. By the end of this book, you will have the skills to determine query performance bottlenecks, avoid pitfalls, and discover the anti-patterns in use.

Foreword by Conor Cunningham, Partner Architect – SQL Server and Azure SQL – Microsoft

What you will learn

- Use Query Store to understand and easily change query performance
- Recognize and eliminate bottlenecks that lead to slow performance
- Deploy quick fixes and long-term solutions to improve query performance
- Implement best practices to minimize performance risk using T-SQL
- Achieve optimal performance by ensuring careful query and index design
- Use the latest performance optimization features in SQL Server 2017 and SQL Server 2019
- Protect query per-

formance during upgrades to newer versions of SQL ServerWho this book is for This book is for database administrators, database developers, data analysts, data scientists, and T-SQL practitioners who want to get started with writing T-SQL code and troubleshooting query performance issues, through the help of practical examples. Previous knowledge of T-SQL querying is not required to get started on this book.

Looks at the principles and clean code, includes case studies showcasing the practices of writing clean code, and contains a list of heuristics and "smells" accumulated from the process of writing clean code.

This book touches on an area seldom explored: the mathematical underpinnings of the relational database. The topic is important, but far too often ignored. This is the first book to explain the underlying math in a way that's accessible to database professionals. Just as importantly, if not more so, this book goes beyond the abstract by showing readers how to apply that math in ways that will make them more productive in their jobs. What's in this book will "open the eyes" of most readers to the great power, elegance, and simplicity inherent in relational database technology.

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.

The demand for SQL information and training continues to grow with the need for a database behind every website capable of offering web-based information queries. SQL is the de facto standard for database retrieval, and if you need to access, update, or

utilize data in a modern database management system, you will need SQL to do it. The Second Edition of Joe Celko's Trees and Hierarchies in SQL for Smarties covers two new sets of extensions over three entirely new chapters and expounds upon the changes that have occurred in SQL standards since the previous edition's publication. Benefit from mastering the challenging aspects of these database applications in SQL as taught by Joe Celko, one of the most-read SQL authors in the world. Expert advice from a noted SQL authority and award-winning columnist who has given 10 years of service to the ANSI SQL standards committee Teaches scores of advanced techniques that can be used with any product, in any SQL environment Offers graph theory and programming techniques for working around deficiencies and gives insight into real-world challenges

"This book takes the somewhat daunting process of database design and breaks it into completely manageable and understandable components. Mike's approach whilst simple is completely professional, and I can recommend this book to any novice database designer." --Sandra Barker, Lecturer, University of South Australia, Australia "Databases are a critical infrastructure technology for information systems and today's business. Mike Hernandez has written a literate explanation of database technology--a topic that is intricate and often obscure. If you design databases yourself, this book will educate you about pitfalls and show you what to do. If you purchase products that use a database, the book explains the technology so that you can understand what the vendor is doing and assess their products better." --Michael Blaha, consultant and trainer, author of A Manager's Guide to Database Technology "If you told me that Mike

Hernandez could improve on the first edition of Database Design for Mere Mortals I wouldn't have believed you, but he did! The second edition is packed with more real-world examples, detailed explanations, and even includes database-design tools on the CD-ROM! This is a must-read for anyone who is even remotely interested in relational database design, from the individual who is called upon occasionally to create a useful tool at work, to the seasoned professional who wants to brush up on the fundamentals. Simply put, if you want to do it right, read this book!" --Matt Greer, Process Control Development, The Dow Chemical Company "Mike's approach to database design is totally common-sense based, yet he's adhered to all the rules of good relational database design. I use Mike's books in my starter database-design class, and I recommend his books to anyone who's interested in learning how to design databases or how to write SQL queries." --Michelle Poolet, President, MVDS, Inc. "Slapping together sophisticated applications with poorly designed data will hurt you just as much now as when Mike wrote his first edition, perhaps even more. Whether you're just getting started developing with data or are a seasoned pro; whether you've read Mike's previous book or this is your first; whether you're happier letting someone else design your data or you love doing it yourself--this is the book for you. Mike's ability to explain these concepts in a way that's not only clear, but fun, continues to amaze me." --From the Foreword by Ken Getz, MCW Technologies, coauthor ASP.NET Developer's JumpStart "The first edition of Mike Hernandez's book Database Design for Mere Mortals was one of the few books that survived the cut when I moved my office to smaller quarters. The second edition expands and improves on the origi-

nal in so many ways. It is not only a good, clear read, but contains a remarkable quantity of clear, concise thinking on a very complex subject. It's a must for anyone interested in the subject of database design." --Malcolm C. Rubel, Performance Dynamics Associates "Mike's excellent guide to relational database design deserves a second edition. His book is an essential tool for fledgling Microsoft Access and other desktop database developers, as well as for client/server pros. I recommend it highly to all my readers." --Roger Jennings, author of Special Edition Using Access 2002 "There are no silver bullets! Database technology has advanced dramatically, the newest crop of database servers perform operations faster than anyone could have imagined six years ago, but none of these technological advances will help fix a bad database design, or capture data that you forgot to include! Database Design for Mere Mortals(TM), Second Edition, helps you design your database right in the first place!" --Matt Nunn, Product Manager, SQL Server, Microsoft Corporation "When my brother started his professional career as a developer, I gave him Mike's book to help him understand database concepts and make real-world application of database technology. When I need a refresher on the finer points of database design, this is the book I pick up. I do not think that there is a better testimony to the value of a book than that it gets used. For this reason I have wholeheartedly recommended to my peers and students that they utilize this book in their day-to-day development tasks." --Chris Kuniciki, Senior Consultant, OfficeZealot.com "Mike has always had an incredible knack for taking the most complex topics, breaking them down, and explaining them so that anyone can 'get it.' He has honed and polished his first very, very good edition and

made it even better. If you're just starting out building database applications, this book is a must-read cover to cover. Expert designers will find Mike's approach fresh and enlightening and a source of great material for training others." --John Viescas, President, Viescas Consulting, Inc., author of *Running Microsoft Access 2000* and coauthor of *SQL Queries for Mere Mortals* "Whether you need to learn about relational database design in general, design a relational database, understand relational database terminology, or learn best practices for implementing a relational database, *Database Design for Mere Mortals(TM), Second Edition*, is an indispensable book that you'll refer to often. With his many years of real-world experience designing relational databases, Michael shows you how to analyze and improve existing databases, implement keys, define table relationships and business rules, and create data views, resulting in data integrity, uniform access to data, and reduced data-entry errors." --Paul Cornell, Site Editor, MSDN Office Developer Center Sound database design can save hours of development time and ensure functionality and reliability. *Database Design for Mere Mortals(TM), Second Edition*, is a straightforward, platform-independent tutorial on the basic principles of relational database design. It provides a commonsense design methodology for developing databases that work. Database design expert Michael J. Hernandez has expanded his best-selling first edition, maintaining its hands-on approach and accessibility while updating its coverage and including even more examples and illustrations. This edition features a CD-ROM that includes diagrams of sample databases, as well as design guidelines, documentation forms, and examples of the database design process. This book will give you the knowledge and tools you

need to create efficient and effective relational databases. Refactoring has proven its value in a wide range of development projects—helping software professionals improve system designs, maintainability, extensibility, and performance. Now, for the first time, leading agile methodologist Scott Ambler and renowned consultant Pramodkumar Sadalage introduce powerful refactoring techniques specifically designed for database systems. Ambler and Sadalage demonstrate how small changes to table structures, data, stored procedures, and triggers can significantly enhance virtually any database design—without changing semantics. You'll learn how to evolve database schemas in step with source code—and become far more effective in projects relying on iterative, agile methodologies. This comprehensive guide and reference helps you overcome the practical obstacles to refactoring real-world databases by covering every fundamental concept underlying database refactoring. Using start-to-finish examples, the authors walk you through refactoring simple standalone database applications as well as sophisticated multi-application scenarios. You'll master every task involved in refactoring database schemas, and discover best practices for deploying refactorings in even the most complex production environments. The second half of this book systematically covers five major categories of database refactorings. You'll learn how to use refactoring to enhance database structure, data quality, and referential integrity; and how to refactor both architectures and methods. This book provides an extensive set of examples built with Oracle and Java and easily adaptable for other languages, such as C#, C++, or VB.NET, and other databases, such as DB2, SQL Server, MySQL, and Sybase. Using this book's techniques and examples, you can

reduce waste, rework, risk, and cost—and build database systems capable of evolving smoothly, far into the future.

The study of software engineering and its applications to system engineering is critical in computer science research. Modern research methodologies, as well as the use of machine and statistical learning in software engineering research, are covered in this book. This book contains the refereed proceedings of the Software Engineering Perspectives in Systems part of the 11th Computer Science On-line Conference 2022 (CSOC 2022), which was held in April 2022 online.

The Complete Guide to Avoiding and Fixing Common Rails 3 Code and Design Problems As developers worldwide have adopted the powerful Ruby on Rails web framework, many have fallen victim to common mistakes that reduce code quality, performance, reliability, stability, scalability, and maintainability. Rails™ AntiPatterns identifies these widespread Rails code and design problems, explains why they're bad and why they happen—and shows exactly what to do instead. The book is organized into concise, modular chapters—each outlines a single common AntiPattern and offers detailed, cookbook-style code solutions that were previously difficult or impossible to find. Leading Rails developers Chad Pytel and Tammer Saleh also offer specific guidance for refactoring existing bad code or design to reflect sound object-oriented principles and established Rails best practices. With their help, developers, architects, and testers can dramatically improve new and existing applications, avoid future problems, and establish superior Rails coding standards throughout their organizations. This book will help you understand, avoid, and solve

problems with Model layer code, from general object-oriented programming violations to complex SQL and excessive redundancy Domain modeling, including schema and database issues such as normalization and serialization View layer tools and conventions Controller-layer code, including RESTful code Service-related APIs, including timeouts, exceptions, backgrounding, and response codes Third-party code, including plug-ins and gems Testing, from test suites to test-driven development processes Scaling and deployment Database issues, including migrations and validations System design for “graceful degradation” in the real world

SQL is full of difficulties and traps for the unwary. You can avoid them if you understand relational theory, but only if you know how to put the theory into practice. In this insightful book, author C.J. Date explains relational theory in depth, and demonstrates through numerous examples and exercises how you can apply it directly to your use of SQL. This second edition includes new material on recursive queries, “missing information” without nulls, new update operators, and topics such as aggregate operators, grouping and ungrouping, and view updating. If you have a modest-to-advanced background in SQL, you'll learn how to deal with a host of common SQL dilemmas. Why is proper column naming so important? Nulls in your database are causing you to get wrong answers. Why? What can you do about it? Is it possible to write an SQL query to find employees who have never been in the same department for more than six months at a time? SQL supports “quantified comparisons,” but they're better avoided. Why? How do you avoid them? Constraints are crucially important, but most SQL products don't support them properly. What

can you do to resolve this situation? Database theory and practice have evolved since the relational model was developed more than 40 years ago. SQL and Relational Theory draws on decades of research to present the most up-to-date treatment of SQL available. C.J. Date has a stature that is unique within the database in-

dustry. A prolific writer well known for the bestselling textbook *An Introduction to Database Systems* (Addison-Wesley), he has an exceptionally clear style when writing about complex principles and theory.