

Database Design For Mere Mortals A Hands On Guide To Relational Database Design

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. Not a reference book, and not a tutorial either, the new second edition of the highly regarded Access Cookbook is an uncommonly useful collection of solutions to problems that Access users and developers are likely to face as they attempt to build increasingly complex applications. Although using any single "recipe" in the book will more than pay back the cost of the book in terms of both hours saved and frustration thwarted, Access Cookbook, Second Edition is much more than a handy assortment of cut-and-paste code. Each of the "recipes" examine a particular problem--problems that commonly occur when you push the upper limits of Access, or ones that are likely to trip up a developer attempting to design a more elegant Access application--even some things you never knew Access could do. The authors then, in a clear, accessible, step-by-step style, present the problems' solution. Following each "recipe" are insights on how Access works, potential pitfalls, interesting programming techniques that are used in the solution, and how and why the solution works, so you can adapt the problem-solving techniques to other similar situations. Fully updated for Access 2003, Access Cookbook, Second Edition is also one of the first books to thoroughly explore new support for .NET managed code and XML. All of the practical, real-world examples have been tested for compatibility with Access 2003, Windows XP, and Windows Server 2003. This updated new edition also covers Access and SharePoint, Access and SmartTags, Access and .NET; and Access and XML. Access power users and programmers at all levels, from the relatively inexperienced to the most sophisticated, will rely on the Access Cookbook for quick solutions to gnarly problems. With a dog-eared copy of Access Cookbook at your side, you can spend your time and energy where it matters most: working on the interesting facets of your

Access application, not just the time-consuming ones.

Fully revised and updated, *Relational Database Design, Second Edition* is the most lucid and effective introduction to relational database design available. Here, you'll find the conceptual and practical information you need to develop a design that ensures data accuracy and user satisfaction while optimizing performance, regardless of your experience level or choice of DBMS.

Supporting the book's step-by-step instruction are three case studies illustrating the planning, analysis, and design steps involved in arriving at a sound design. These real-world examples include object-relational design techniques, which are addressed in greater detail in a new chapter devoted entirely to this timely subject. * Concepts you need to master to put the book's practical instruction to work. * Methods for tailoring your design to the environment in which the database will run and the uses to which it will be put. * Design approaches that ensure data accuracy and consistency. * Examples of how design can inhibit or boost database application performance. * Object-relational design techniques, benefits, and examples. * Instructions on how to choose and use a normalization technique. * Guidelines for understanding and applying Codd's rules. * Tools to implement a relational design using SQL. * Techniques for using CASE tools for database design.

Congratulations! You completed the MongoDB application within the given tight timeframe and there is a party to celebrate your application's release into production. Although people are congratulating you at the celebration, you are feeling some uneasiness inside. To complete the project on time required making a lot of assumptions about the data, such as what terms meant and how calculations are derived. In addition, the poor documentation about the application will be of limited use to the support team, and not investigating all of the inherent rules in the data may eventually lead to poorly-performing structures in the not-so-distant future. Now, what if you had a time machine and could go back and read this book. You would learn that even NoSQL databases like MongoDB require some level of data modeling. Data modeling is the process of learning about the data, and regardless of technology, this process must be performed for a successful application. You would learn the value of conceptual, logical, and physical data modeling and how each stage increases our knowledge of the data and reduces assumptions and poor design decisions. Read this book to learn how to do data modeling for MongoDB applications, and accomplish these five objectives: Understand how data modeling contributes to the process of learning about the data, and is, therefore, a required technique, even when the resulting database is not relational. That is, NoSQL does not mean NoDataModeling! Know how NoSQL databases differ from traditional relational databases, and where MongoDB fits. Explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts, and learn the basics of adding, querying, updating, and deleting data in MongoDB. Practice a streamlined, template-driven approach to performing conceptual, logical, and physical data modeling. Recognize that data modeling does not always have to lead to traditional data models! Distinguish top-down from bottom-up development approaches and complete a top-down case study which ties all of the modeling techniques together. This book is written for anyone who is working with, or will be working with MongoDB, including business analysts, data modelers, database administrators, developers, project managers, and data scientists. There are three sections: In Section I,

Getting Started, we will reveal the power of data modeling and the tight connections to data models that exist when designing any type of database (Chapter 1), compare NoSQL with traditional relational databases and where MongoDB fits (Chapter 2), explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts (Chapter 3), and explain the basics of adding, querying, updating, and deleting data in MongoDB (Chapter 4). In Section II, Levels of Granularity, we cover Conceptual Data Modeling (Chapter 5), Logical Data Modeling (Chapter 6), and Physical Data Modeling (Chapter 7). Notice the “ing” at the end of each of these chapters. We focus on the process of building each of these models, which is where we gain essential business knowledge. In Section III, Case Study, we will explain both top down and bottom up development approaches and go through a top down case study where we start with business requirements and end with the MongoDB database. This case study will tie together all of the techniques in the previous seven chapters. Nike Senior Data Architect Ryan Smith wrote the foreword. Key points are included at the end of each chapter as a way to reinforce concepts. In addition, this book is loaded with hands-on exercises, along with their answers provided in Appendix A. Appendix B contains all of the book’s references and Appendix C contains a glossary of the terms used throughout the text.

Through clear language, step-by-step discussions, and quizzes at the end of each chapter, the author makes databases easy. Quickly learn the core skills needed to design, configure, manage, and manipulate databases, whether at work or at home. Topics such as exploring different database models, planning their design, minimizing redundant data, designing tables, applying database design concepts, and implementing database security are covered. This is that fast, easy-to-understand tutorial that you’ve been looking for.

&a>Database Design for Mere Mortals™, 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 book will give you the knowledge and tools you need to create efficient and effective relational databases.

This block is concerned with the database lifecycle, which describes the stages a database goes through, from the time the need for a database is established until it is withdrawn from use. This block applies the practice developed in Block 3 to systematically develop, implement and maintain a database design that supports the information requirements of an enterprise. It presents a simple framework for database development and maintenance. This is a very practical block and will require you to write and execute SQL statements for which you will need access to a computer installed with the course software (order code M359/CDR01) and database cards Scenarios and Hospital conceptual data model (order code M359/DBCARDS)

Algorithms play an important role in both the science and practice of computing. To optimally use algorithms, a deeper understanding of their logic and mathematics is essential. Beyond traditional computing, the ability to apply these algorithms to solve real-world problems is a necessary skill, and this is what this book focuses on.

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.

The aim of this work is to provide a correct and up-to-date understanding of the practical aspects of crucial, yet little-understood core database issues. The author identifies fundamental concepts, principles, and techniques and assesses the treatment of those issues in SQL (both the standard and commercial implementations) and gives advice on how to deal with them. Topics covered include complex data types, missing information, data hierarchies, and quota queries.

Annotation copyrighted by Book News, Inc., Portland, OR

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.

"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 original 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 Kunicki, 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.

Build Android 6 Material Design Apps That Are Stunningly Attractive, Functional, and Intuitive As Android development has matured and grown increasingly competitive, developers have recognized the crucial importance of good design. With Material Design, Google introduced its most radical visual changes ever, and made effective design even more essential. Android 6 and the design support library continue to push mobile design forward. In Android User Interface Design, Second Edition, leading Android developer and user experience (UX) advocate Ian G. Clifton shows how to combine exceptional usability and outstanding visual appeal. Clifton helps you build apps that new users can succeed with instantly: apps that leverage users' previous experience, reflect platform conventions, and

never test their patience. You won't need any design experience: Clifton walks you through the entire process, from wireframes and flowcharts to finished apps with polished animations and advanced compositing. You'll find hands-on case studies and extensive downloadable sample code, including complete finished apps.

- Integrate Material Design into backward compatible Android 6 apps
- Understand views, the building blocks of Android user interfaces
- Make the most of wireframes and conceptual prototypes
- Apply user-centered design throughout
- Master the essentials of typography and iconography
- Use custom themes and styles for consistent visuals
- Handle inputs and scrolling
- Create beautiful transition animations
- Use advanced components like spans and image caches
- Work with the canvas, color filters, shaders, and image compositing
- Combine multiple views into efficient custom components
- Customize views to meet unique drawing or interaction requirements
- Maximize downloads by designing compelling app store assets

Step by step, this guide bridges the gap between Android developers and designers, so you can collaborate on world-class app designs...or do it all yourself! "This well-presented, easy-to-grasp book gets to the heart of Android User Interface Design. Well worth the reading time!" --Dr. Adam Porter, University of Maryland, Fraunhofer Center for Experimental Software Engineering "Ian's grasp of Android is fantastic, and this book is a great read for any developer or designer. I've personally worked on 30+ Android applications, and I was learning new tips with every chapter." --Cameron Banga, Lead Designer, 9magnets, LLC

Database design expert Michael J. Hernandez introduces the core concepts of design theory and method without the technical jargon. With its hands-on approach and a wealth of practical examples, Database Design for Mere Mortals will provide any developer with a common-sense design methodology for developing databases that work.

Summary Big Data teaches you to build big data systems using an architecture that takes advantage of clustered hardware along with new tools designed specifically to capture and analyze web-scale data. It describes a scalable, easy-to-understand approach to big data systems that can be built and run by a small team. Following a realistic example, this book guides readers through the theory of big data systems, how to implement them in practice, and how to deploy and operate them once they're built. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book Web-scale applications like social networks, real-time analytics, or e-commerce sites deal with a lot of data, whose volume and velocity exceed the limits of traditional database systems. These applications require architectures built around clusters of machines to store and process data of any size, or speed. Fortunately, scale and simplicity are not mutually exclusive. Big Data teaches you to build big data systems using an architecture designed specifically to capture and analyze web-scale data. This book presents the Lambda Architecture, a scalable, easy-to-understand approach that can be built and run by a small team. You'll explore the theory of big data

systems and how to implement them in practice. In addition to discovering a general framework for processing big data, you'll learn specific technologies like Hadoop, Storm, and NoSQL databases. This book requires no previous exposure to large-scale data analysis or NoSQL tools. Familiarity with traditional databases is helpful. What's Inside Introduction to big data systems Real-time processing of web-scale data Tools like Hadoop, Cassandra, and Storm Extensions to traditional database skills About the Authors Nathan Marz is the creator of Apache Storm and the originator of the Lambda Architecture for big data systems. James Warren is an analytics architect with a background in machine learning and scientific computing. Table of Contents A new paradigm for Big Data PART 1 BATCH LAYER Data model for Big Data Data model for Big Data: Illustration Data storage on the batch layer Data storage on the batch layer: Illustration Batch layer Batch layer: Illustration An example batch layer: Architecture and algorithms An example batch layer: Implementation PART 2 SERVING LAYER Serving layer Serving layer: Illustration PART 3 SPEED LAYER Realtime views Realtime views: Illustration Queuing and stream processing Queuing and stream processing: Illustration Micro-batch stream processing Micro-batch stream processing: Illustration Lambda Architecture in depth 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. Fully updated and expanded from the previous edition, *A Practical Guide to Database Design, Second Edition*, is intended for those involved in the design or development of a database system or application. It begins by focusing on how to create a logical data model where data is stored "where it belongs." Next, data usage is reviewed to transform the logical model into a physical data model that will satisfy user performance requirements. Finally, it describes how to use various software tools to create user interfaces to review and update data in a database. Organized into 11 chapters, the book begins with an overview of the functionality of database management systems and how they guarantee the accuracy and availability of data. It then describes how to define and normalize data requirements to create a logical data model, then map them into an initial solution for a physical database. The book next presents how to use an industry-leading data modeling tool to define and manage logical and physical data models. After that, it describes how to implement a physical database using either Microsoft Access or SQL Server and how to use Microsoft Access to create windows interfaces to query or update data in tables. The last part of the book reviews software tools and explores the design and implementation of a database using as an example a much more complex data environment for a University. The book ends with a description of how to use PHP to build a web-based interface to review and update data in a database.

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.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. Hands-On Database uses a scenario-based approach that shows readers how to build a database by providing them with the context of a running case throughout each step of the process.

"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.

*The fastest, easiest way to master SQL-no experience required *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. Now you can learn SQL hands on--and get the business information you need, faster than you ever imagined This package brings together the industry's #1 interactive SQL training CD-ROM, the SQL Queries Multimedia Cyber Classroom, with the world's best SQL book for business professionals, SQL Queries for Mere Mortals No matter what database you use, you'll master SQL by practicing with hundreds of live queries and exercises--and learning from hours of animations and audio commentary by two world-renowned database consultants *Over 200 animated figures, with expert audio commentary, illustrate key concepts *Learning SQL is just one click away searchable version of SQL Queries for Mere Morals *Test your knowledge with 100+ self-review questions and 100+ practice exercises *INCLUDES THE #1 BUSINESS GUIDE TO SQL *BONUS: SECOND CD-ROM contains five real-world sample databases designed to support virtually any practice query. *Starts from scratch No unnecessary jargon: just plain-English help and useful examples *Real-world focus: Solutions for the problems businesspeople really encounter. *Up to date Covers the latest ANSI standard SQL--skills apply to virtually any current desktop or enterprise database *Includes advanced, high-performance techniques for solving even complex SQL problems. Queries Training Course covers everything you need to know to succeed with SQL, including concepts, syntax, troubleshooting techniques, performance optimization, and more: *Basic relational database concepts *Fundamental SQL concepts *SELECT statements *Creating expressions *Applying filters *Working with joins *Integrating data from multiple tables *Grouping data *Totaling data *Summarizing data *Updating data *Inserting data *Deleting data *And much more -- including advanced techniques Requirements Windows(r) 95/98/2000/NT(r) 4.0 Internet Explorer (included) 20 MB disk space 32 MB RAM CD-ROM drive Sound card support

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

NoSQL was developed to overcome the limitations of relational databases in the largest Web applications at companies such as Google, Yahoo and Facebook. As it is applied more widely, developers are finding that it can simplify scalability while requiring far less coding and management overhead. However, NoSQL requires fundamentally different approaches to database design and modeling, and many conventional relational techniques lead to suboptimal results. *NoSQL for Mere Mortals* is an easy, practical guide to succeeding with NoSQL in your environment. Following the classic, best-selling format pioneered in *SQL Queries for Mere Mortals*, enterprise database expert Dan Sullivan guides you step-by-step through choosing technologies, designing high-performance databases, and planning for long-term maintenance. *Sullivan* introduces each type of NoSQL database, shows how to install and manage them, and demonstrates how to leverage their features while avoiding common mistakes that lead to poor performance and unmet requirements. He uses four popular NoSQL databases as reference models: MongoDB, a document database; Cassandra, a column family data store; Redis, a key-value database; and Neo4j, a graph database. You'll find explanations of each database's structure and capabilities, practical guidelines for choosing amongst them, and expert guidance on designing databases with them. *Packed with examples, NoSQL for Mere Mortals is today's best way to master NoSQL -- whether you're a DBA, developer, user, or student.*

In just 24 lessons of one hour or less, you will learn professional techniques to design and build efficient databases and query them to extract useful information. Using a straightforward, step-by-step approach, each lesson builds on the previous one, allowing you to learn the essentials of ANSI SQL from the ground up. Example code demonstrates the

authors' professional techniques, while exercises written for MySQL offer the reader hands-on learning with an open-source database. Included are advanced techniques for using views, managing transactions, database administration, and extending SQL. Step-by-step instructions carefully walk you through the most common SQL tasks. Q&As, Quizzes, and Exercises at the end of each chapter help you test your knowledge. Notes and Tips point out shortcuts and solutions. New terms are clearly defined and explained. Learn how to... Use SQL-2003, the latest standard for the Structured Query Language Design and deploy efficient, secure databases Build advanced queries for information retrieval Sort, group, and summarize information for best presentation Tune databases and queries for maximum performance Understand database administration and security techniques For more than ten years the authors have studied, applied, and documented the SQL standard and its application to critical database systems. Ryan Stephens and Ron Plew are entrepreneurs, speakers, and cofounders of Perpetual Technologies, Inc. (PTI), a fast-growing IT management and consulting firm which specializes in database technologies. They taught database courses for Indiana University–Purdue University in Indianapolis for five years and have authored more than a dozen books on Oracle, SQL, database design, and the high availability of critical systems. Arie D. Jones is Senior SQL Server database administrator and analyst for PTI. He is a regular speaker at technical events and has authored several books and articles. Category: Database Covers: ANSI SQL User Level: Beginning–Intermediate Register your book at informit.com/title/9780672330186 for convenient access to updates and corrections as they become available.

The #1 Easy, Commonsense Guide to Database Design Now Updated Foreword by Michelle Poollet, Mount Vernon Data Systems LLC Michael J. Hernandezs best-selling Database Design for Mere Mortals has earned worldwide respect as the simplest way to learn relational database design. Now, hes made this hands-on, software independent tutorial even clearer and easier to use. Step by step, this new 25th Anniversary Edition shows you how to design modern databases that are soundly structured, reliable, and flexible, even in the latest online applications. Hernandez guides you through everything from planning to defining tables, fields, keys, table relationships, business rules, and views. Youll learn practical ways to improve data integrity, how to avoid common mistakes, and when to break the rules. Updated review questions and figures help you learn these techniques more easily and effectively. Understand database types, models, and design terminology Perform interviews to efficiently capture requirements even if everyones working remotely Set clear design objectives and transform them into effective designs Analyze a current database so you can identify ways to improve it Establish table structures and relationships, assign primary keys, set field specifications, and set up views Ensure the correct level of data integrity for each database Identify and establish business rules Preview and prepare for the future of relational databases Whatever relational database systems you use, Hernandez will help you design

databases that are robust and trustworthy. Never designed a database before? Settling for inadequate generic designs? Running existing databases that need improvement? Start here.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. The Language of SQL, Second Edition Many SQL texts attempt to serve as an encyclopedic reference on SQL syntax -- an approach that is often counterproductive, because that information is readily available in online references published by the major database vendors. For SQL beginners, it's more important for a book to focus on general concepts and to offer clear explanations and examples of what various SQL statements can accomplish. This is that book. A number of features make The Language of SQL unique among introductory SQL books. First, you will not be required to download software or sit with a computer as you read the text. The intent of this book is to provide examples of SQL usage that can be understood simply by reading. Second, topics are organized in an intuitive and logical sequence. SQL keywords are introduced one at a time, allowing you to grow your understanding as you encounter new terms and concepts. Finally, this book covers the syntax of three widely used databases: Microsoft SQL Server, MySQL, and Oracle. Special "Database Differences" sidebars clearly show you any differences in syntax among these three databases, and instructions are included on how to obtain and install free versions of the databases. This is the only book you need to gain a quick working knowledge of SQL and relational databases. ·Learn How To... Use SQL to retrieve data from relational databases Apply functions and calculations to data Group and summarize data in a variety of useful ways Use complex logic to retrieve only the data you need Update data and create new tables Design relational databases so that data retrieval is easy and intuitive Use spreadsheets to transform your data into meaningful displays Retrieve data from multiple tables via joins, subqueries, views, and set logic Create, modify, and execute stored procedures Install Microsoft SQL Server, MySQL, or Oracle

Fully revised, updated, and expanded, Relational Database Design and Implementation, Third Edition is the most lucid and effective introduction to the subject available for IT/IS professionals interested in honing their skills in database design, implementation, and administration. This book provides the conceptual and practical information necessary to develop a design and management scheme that ensures data accuracy and user satisfaction while optimizing performance, regardless of experience level or choice of DBMS. The book begins by reviewing basic concepts of databases and database design, then briefly reviews the SQL one would use to create databases. Topics such as the relational data model, normalization, data entities and Codd's Rules (and why they are important) are covered clearly and concisely but without resorting to "Dummies"-style talking down to the reader. Supporting the book's step-by-step instruction are three NEW case studies illustrating database planning, analysis, design, and management practices. In

addition to these real-world examples, which include object-relational design techniques, an entirely NEW section consisting of three chapters is devoted to database implementation and management issues. * Principles needed to understand the basis of good relational database design and implementation practices. * Examples to illustrate core concepts for enhanced comprehension and to put the book's practical instruction to work. * Methods for tailoring DB design to the environment in which the database will run and the uses to which it will be put. * Design approaches that ensure data accuracy and consistency. * Examples of how design can inhibit or boost database application performance. * Object-relational design techniques, benefits, and examples. * Instructions on how to choose and use a normalization technique. * Guidelines for understanding and applying Codd's rules. * Tools to implement a relational design using SQL. * Techniques for using CASE tools for database design.

Learn effective and scalable database design techniques in a SQL Server environment. Pro SQL Server 2012 Relational Database Design and Implementation covers everything from design logic that business users will understand, all the way to the physical implementation of design in a SQL Server database. Grounded in best practices and a solid understanding of the underlying theory, Louis Davidson shows how to “get it right” in SQL Server database design and lay a solid groundwork for the future use of valuable business data. Gives a solid foundation in best practices and relational theory Covers the latest implementation features in SQL Server Takes you from conceptual design to an effective, physical implementation

The #1 Easy, Commonsense Guide to Database Design! Michael J. Hernandez’s best-selling Database Design for Mere Mortals® has earned worldwide respect as the clearest, simplest way to learn relational database design. Now, he’s made this hands-on, software-independent tutorial even easier, while ensuring that his design methodology is still relevant to the latest databases, applications, and best practices. Step by step, Database Design for Mere Mortals®, Third Edition, shows you how to design databases that are soundly structured, reliable, and flexible, even in modern web applications. Hernandez guides you through everything from database planning to defining tables, fields, keys, table relationships, business rules, and views. You’ll learn practical ways to improve data integrity, how to avoid common mistakes, and when to break the rules. Coverage includes Understanding database types, models, and design terminology Discovering what good database design can do for you—and why bad design can make your life miserable Setting objectives for your database, and transforming those objectives into real designs Analyzing a current database so you can identify ways to improve it Establishing table structures and relationships, assigning primary keys, setting field specifications, and setting up views Ensuring the appropriate level of data integrity for each application Identifying and establishing business rules Whatever relational database systems you use, Hernandez will help you design databases

that are robust and trustworthy. Never designed a database before? Settling for inadequate generic designs? Running existing databases that need improvement? Start here.

The latest edition of a popular text and reference on database research, with substantial new material and revision; covers classical literature and recent hot topics. Lessons from database research have been applied in academic fields ranging from bioinformatics to next-generation Internet architecture and in industrial uses including Web-based e-commerce and search engines. The core ideas in the field have become increasingly influential. This text provides both students and professionals with a grounding in database research and a technical context for understanding recent innovations in the field. The readings included treat the most important issues in the database area--the basic material for any DBMS professional. This fourth edition has been substantially updated and revised, with 21 of the 48 papers new to the edition, four of them published for the first time. Many of the sections have been newly organized, and each section includes a new or substantially revised introduction that discusses the context, motivation, and controversies in a particular area, placing it in the broader perspective of database research. Two introductory articles, never before published, provide an organized, current introduction to basic knowledge of the field; one discusses the history of data models and query languages and the other offers an architectural overview of a database system. The remaining articles range from the classical literature on database research to treatments of current hot topics, including a paper on search engine architecture and a paper on application servers, both written expressly for this edition. The result is a collection of papers that are seminal and also accessible to a reader who has a basic familiarity with database systems.

A guide to relational database design covers such topics as setting objectives, establishing table structures, and identifying and establishing business rules.

Learn effective and scalable database design techniques in SQL Server 2019 and other recent SQL Server versions. This book is revised to cover additions to SQL Server that include SQL graph enhancements, in-memory online transaction processing, temporal data storage, row-level security, and other design-related features. This book will help you design OLTP databases that are high-quality, protect the integrity of your data, and perform fast on-premises, in the cloud, or in hybrid configurations. Designing an effective and scalable database using SQL Server is a task requiring skills that have been around for well over 30 years, using technology that is constantly changing. This book covers everything from design logic that business users will understand to the physical implementation of design in a SQL Server database. Grounded in best practices and a solid understanding of the underlying theory, author Louis Davidson shows you how to "get it right" in SQL Server database design and lay a solid groundwork for the future use of valuable business data. What You Will Learn Develop conceptual models of client data using interviews and client documentation

Implement designs that work on premises, in the cloud, or in a hybrid approach Recognize and apply common database design patterns Normalize data models to enhance integrity and scalability of your databases for the long-term use of valuable data Translate conceptual models into high-performing SQL Server databases Secure and protect data integrity as part of meeting regulatory requirements Create effective indexing to speed query performance Understand the concepts of concurrency Who This Book Is For Programmers and database administrators of all types who want to use SQL Server to store transactional data. The book is especially useful to those wanting to learn the latest database design features in SQL Server 2019 (features that include graph objects, in-memory OLTP, temporal data support, and more). Chapters on fundamental concepts, the language of database modeling, SQL implementation, and the normalization process lay a solid groundwork for readers who are just entering the field of database design. More advanced chapters serve the seasoned veteran by tackling the latest in physical implementation features that SQL Server has to offer. The book has been carefully revised to cover all the design-related features that are new in SQL Server 2019.

Database Design for Mere Mortals A Hands-on Guide to Relational Database Design Pearson Education

Data is getting bigger and more complex by the day, and so are your choices in handling it. Explore some of the most cutting-edge databases available - from a traditional relational database to newer NoSQL approaches - and make informed decisions about challenging data storage problems. This is the only comprehensive guide to the world of NoSQL databases, with in-depth practical and conceptual introductions to seven different technologies: Redis, Neo4J, CouchDB, MongoDB, HBase, Postgres, and DynamoDB. This second edition includes a new chapter on DynamoDB and updated content for each chapter. While relational databases such as MySQL remain as relevant as ever, the alternative, NoSQL paradigm has opened up new horizons in performance and scalability and changed the way we approach data-centric problems. This book presents the essential concepts behind each database alongside hands-on examples that make each technology come alive. With each database, tackle a real-world problem that highlights the concepts and features that make it shine. Along the way, explore five database models - relational, key/value, columnar, document, and graph - from the perspective of challenges faced by real applications. Learn how MongoDB and CouchDB are strikingly different, make your applications faster with Redis and more connected with Neo4J, build a cluster of HBase servers using cloud services such as Amazon's Elastic MapReduce, and more. This new edition brings a brand new chapter on DynamoDB, updated code samples and exercises, and a more up-to-date account of each database's feature set. Whether you're a programmer building the next big thing, a data scientist seeking solutions to thorny problems, or a technology enthusiast venturing into new territory, you will find something to inspire you in this book. What You Need: You'll need a *nix shell (Mac OS or Linux preferred, Windows users will need Cygwin), Java 6 (or greater), and Ruby 1.8.7 (or greater). Each chapter will list the downloads required for that database.

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.

The purpose of *ASP.NET JumpStart* is to show readers the practical applications of .NET and ASP.NET by illustrating how to build Web-based applications using Web Forms and Web Services. Emphasis will be on good programming standards and practices. The reader will be taken from an introduction of the VB .NET language to intermediate topics through a step-by-step approach, which gives the reader the opportunity to try out the practices presented in each chapter.

Salary surveys worldwide regularly place software architect in the top 10 best jobs, yet no real guide exists to help developers become architects. Until now. This book provides the first comprehensive overview of software architecture's many aspects. Aspiring and existing architects alike will examine architectural characteristics, architectural patterns, component determination, diagramming and presenting architecture, evolutionary architecture, and many other topics. Mark Richards and Neal Ford—hands-on practitioners who have taught software architecture classes professionally for years—focus on architecture principles that apply across all technology stacks. You'll explore software architecture in a modern light, taking into account all the innovations of the past decade. This book examines:

- Architecture patterns: The technical basis for many architectural decisions
- Components: Identification, coupling, cohesion, partitioning, and granularity
- Soft skills: Effective team management, meetings, negotiation, presentations, and more
- Modernity: Engineering practices and operational approaches that have changed radically in the past few years
- Architecture as an engineering discipline: Repeatable results, metrics, and concrete valuations that add rigor to software architecture

The rapidly increasing volume of information contained in relational databases places a strain on databases, performance, and maintainability: DBAs are under greater pressure than ever to optimize database structure for system performance and administration. *Physical Database Design* discusses the concept of how physical structures of databases affect performance, including specific examples, guidelines, and best and worst practices for a variety of DBMSs and configurations. Something as simple as improving the table index design has a profound impact on performance. Every form of relational database, such as Online Transaction Processing (OLTP), Enterprise Resource Management (ERP), Data Mining (DM), or Management Resource Planning (MRP), can be improved using the methods provided in the book. The first complete treatment on physical database

design, written by the authors of the seminal, Database Modeling and Design: Logical Design, Fourth Edition Includes an introduction to the major concepts of physical database design as well as detailed examples, using methodologies and tools most popular for relational databases today: Oracle, DB2 (IBM), and SQL Server (Microsoft) Focuses on physical database design for exploiting B+tree indexing, clustered indexes, multidimensional clustering (MDC), range partitioning, shared nothing partitioning, shared disk data placement, materialized views, bitmap indexes, automated design tools, and more!

[Copyright: 77b423c24808b389bfe9a650f3aed0d4](#)