

Greenplum Interview Questions And Answers Part1 Dwgeek Com

Clear your doubts about Business Intelligence and start your new journey KEY FEATURES ? Includes successful methods and innovative ideas to achieve success with BI. ? Vendor-neutral, unbiased, and based on experience. ? Highlights practical challenges in BI journeys. ? Covers financial aspects along with technical aspects. ? Showcases multiple BI organization models and the structure of BI teams. DESCRIPTION The book demystifies misconceptions and misinformation about BI. It provides clarity to almost everything related to BI in a simplified and unbiased way. It covers topics right from the definition of BI, terms used in the BI definition, coinage of BI, details of the different main uses of BI, processes that support the main uses, side benefits, and the level of importance of BI, various types of BI based on various parameters, main phases in the BI journey and the challenges faced in each of the phases in the BI journey. It clarifies myths about self-service BI and real-time BI. The book covers the structure of a typical internal BI team, BI organizational models, and the main roles in BI. It also clarifies the doubts around roles in BI. It explores the different components that add to the cost of BI and explains how to calculate the total cost of the ownership of BI and ROI for BI. It covers several ideas, including unconventional ideas to achieve BI success and also learn about IBI. It explains the different types of BI architectures, commonly used technologies, tools, and concepts in BI and provides clarity about the boundary of BI w.r.t technologies, tools, and concepts. The book helps you lay a very strong foundation and provides the right perspective about BI. It enables you to start or restart your journey with BI. WHAT YOU WILL LEARN ? Builds a strong conceptual foundation in BI. ? Gives the right perspective and clarity on BI uses, challenges, and architectures. ? Enables you to make the right decisions on the BI structure, organization model, and budget. ? Explains which type of BI solution is required for your business. ? Applies successful BI ideas. WHO THIS BOOK IS FOR This book is a must-read for business managers, BI aspirants, CxOs, and all those who want to drive the business value with data-driven insights. TABLE OF CONTENTS 1. What is Business Intelligence? 2. Why do Businesses need BI? 3. Types of Business Intelligence 4. Challenges in Business Intelligence 5. Roles in Business Intelligence 6. Financials of Business Intelligence 7. Ideas for Success with BI 8. Introduction to IBI 9. BI Architectures 10. Demystify Tech, Tools, and Concepts in BI

Big Data Imperatives, focuses on resolving the key questions on everyone's mind: Which data matters? Do you have enough data volume to justify the usage? How you want to process this amount of data? How long do you really need to keep it active for your analysis, marketing, and BI applications? Big data is emerging from the realm of one-off projects to mainstream business adoption; however, the real value of big data is not in the overwhelming size of it, but more in its effective use. This book addresses the following big data characteristics: Very large, distributed aggregations of loosely structured data – often incomplete and inaccessible Petabytes/Exabytes of data Millions/billions of people providing/contributing to the context behind the data Flat schema's with few complex interrelationships Involves time-stamped events Made up of incomplete data Includes connections between data elements that must be probabilistically inferred Big Data Imperatives explains 'what big data can do'. It can batch

process millions and billions of records both unstructured and structured much faster and cheaper. Big data analytics provide a platform to merge all analysis which enables data analysis to be more accurate, well-rounded, reliable and focused on a specific business capability. Big Data Imperatives describes the complementary nature of traditional data warehouses and big-data analytics platforms and how they feed each other. This book aims to bring the big data and analytics realms together with a greater focus on architectures that leverage the scale and power of big data and the ability to integrate and apply analytics principles to data which earlier was not accessible. This book can also be used as a handbook for practitioners; helping them on methodology, technical architecture, analytics techniques and best practices. At the same time, this book intends to hold the interest of those new to big data and analytics by giving them a deep insight into the realm of big data.

Cowritten by Ralph Kimball, the world's leading data warehousing authority, whose previous books have sold more than 150,000 copies Delivers real-world solutions for the most time- and labor-intensive portion of data warehousing-data staging, or the extract, transform, load (ETL) process Delineates best practices for extracting data from scattered sources, removing redundant and inaccurate data, transforming the remaining data into correctly formatted data structures, and then loading the end product into the data warehouse Offers proven time-saving ETL techniques, comprehensive guidance on building dimensional structures, and crucial advice on ensuring data quality

Joe Celko's SQL Puzzles and Answers, Second Edition, challenges you with his trickiest puzzles and then helps solve them with a variety of solutions and explanations. Author Joe Celko demonstrates the thought processes that are involved in attacking a problem from an SQL perspective to help advanced database programmers solve the puzzles you frequently face. These techniques not only help with the puzzle at hand, but also help develop the mindset needed to solve the many difficult SQL puzzles you face every day. This updated edition features many new puzzles; dozens of new solutions to puzzles; and new chapters on temporal query puzzles and common misconceptions about SQL and RDBMS that leads to problems. This book is recommended for database programmers with a good knowledge of SQL. A great collection of tricky SQL puzzles with a variety of solutions and explanations Uses the proven format of puzzles and solutions to provide a user-friendly, practical look into SQL programming problems - many of which will help users solve their own problems New edition features: Many new puzzles added!, Dozens of new solutions to puzzles, and using features in SQL-99, Code is edited to conform to SQL STYLE rules, New chapter on temporal query puzzles, New chapter on common misconceptions about SQL and RDBMS that leads to problems

Analyze vast amounts of data in record time using Apache Spark with Databricks in the Cloud. Learn the fundamentals, and more, of running analytics on large clusters in Azure and AWS, using Apache Spark with Databricks on top. Discover how to squeeze the most value out of your data at a mere fraction of what classical analytics solutions cost, while at the same time getting the results you need, incrementally faster. This book explains how the confluence of these pivotal technologies gives you enormous power, and cheaply, when it comes to huge datasets. You will begin by learning how cloud infrastructure makes it possible to scale your code to large amounts of processing units, without having to pay for the machinery in advance. From there you will learn how

Apache Spark, an open source framework, can enable all those CPUs for data analytics use. Finally, you will see how services such as Databricks provide the power of Apache Spark, without you having to know anything about configuring hardware or software. By removing the need for expensive experts and hardware, your resources can instead be allocated to actually finding business value in the data. This book guides you through some advanced topics such as analytics in the cloud, data lakes, data ingestion, architecture, machine learning, and tools, including Apache Spark, Apache Hadoop, Apache Hive, Python, and SQL. Valuable exercises help reinforce what you have learned. What You Will Learn Discover the value of big data analytics that leverage the power of the cloud Get started with Databricks using SQL and Python in either Microsoft Azure or AWS Understand the underlying technology, and how the cloud and Apache Spark fit into the bigger picture See how these tools are used in the real world Run basic analytics, including machine learning, on billions of rows at a fraction of a cost or free Who This Book Is For Data engineers, data scientists, and cloud architects who want or need to run advanced analytics in the cloud. It is assumed that the reader has data experience, but perhaps minimal exposure to Apache Spark and Azure Databricks. The book is also recommended for people who want to get started in the analytics field, as it provides a strong foundation.

Data Science and Big Data Analytics is about harnessing the power of data for new insights. The book covers the breadth of activities and methods and tools that Data Scientists use. The content focuses on concepts, principles and practical applications that are applicable to any industry and technology environment, and the learning is supported and explained with examples that you can replicate using open-source software. This book will help you: Become a contributor on a data science team Deploy a structured lifecycle approach to data analytics problems Apply appropriate analytic techniques and tools to analyzing big data Learn how to tell a compelling story with data to drive business action Prepare for EMC Proven Professional Data Science Certification Corresponding data sets are available from the book's page at Wiley which you can find on the Wiley site by searching for the ISBN 9781118876138. Get started discovering, analyzing, visualizing, and presenting data in a meaningful way today!

Re-architect relational applications to NoSQL, integrate relational database management systems with the Hadoop ecosystem, and transform and migrate relational data to and from Hadoop components. This book covers the best-practice design approaches to re-architecting your relational applications and transforming your relational data to optimize concurrency, security, denormalization, and performance. Winner of IBM's 2012 Gerstner Award for his implementation of big data and data warehouse initiatives and author of Practical Hadoop Security, author Bhushan Lakhe walks you through the entire transition process. First, he lays out the criteria for deciding what blend of re-architecting, migration, and integration between RDBMS and HDFS best meets your transition objectives. Then he demonstrates how to design your transition model. Lakhe proceeds to cover the selection criteria for ETL tools, the implementation steps for migration with SQOOP- and Flume-based data transfers, and transition optimization techniques for tuning partitions, scheduling aggregations, and redesigning ETL. Finally, he assesses the pros and cons of data lakes and Lambda architecture as integrative solutions and illustrates their implementation with real-world case

studies. Hadoop/NoSQL solutions do not offer by default certain relational technology features such as role-based access control, locking for concurrent updates, and various tools for measuring and enhancing performance. Practical Hadoop Migration shows how to use open-source tools to emulate such relational functionalities in Hadoop ecosystem components. What You'll Learn Decide whether you should migrate your relational applications to big data technologies or integrate them Transition your relational applications to Hadoop/NoSQL platforms in terms of logical design and physical implementation Discover RDBMS-to-HDFS integration, data transformation, and optimization techniques Consider when to use Lambda architecture and data lake solutions Select and implement Hadoop-based components and applications to speed transition, optimize integrated performance, and emulate relational functionalities Who This Book Is For Database developers, database administrators, enterprise architects, Hadoop/NoSQL developers, and IT leaders. Its secondary readership is project and program managers and advanced students of database and management information systems.

A century ago, daily life ground to a halt when the circus rolled into town. Across America, banks closed, schools canceled classes, farmers left their fields, and factories shut down so that everyone could go to the show. In this entertaining and provocative book, Janet Davis links the flowering of the early-twentieth-century American railroad circus to such broader historical developments as the rise of big business, the breakdown of separate spheres for men and women, and the genesis of the United States' overseas empire. In the process, she casts the circus as a powerful force in consolidating the nation's identity as a modern industrial society and world power. Davis explores the multiple "shows" that took place under the big top, from scripted performances to exhibitions of laborers assembling and tearing down tents to impromptu spectacles of audiences brawling, acrobats falling, and animals rampaging. Turning Victorian notions of gender, race, and nationhood topsy-turvy, the circus brought its vision of a rapidly changing world to spectators--rural as well as urban--across the nation. Even today, Davis contends, the influence of the circus continues to resonate in popular representations of gender, race, and the wider world.

Ready to unlock the power of your data? With this comprehensive guide, you'll learn how to build and maintain reliable, scalable, distributed systems with Apache Hadoop. This book is ideal for programmers looking to analyze datasets of any size, and for administrators who want to set up and run Hadoop clusters. You'll find illuminating case studies that demonstrate how Hadoop is used to solve specific problems. This third edition covers recent changes to Hadoop, including material on the new MapReduce API, as well as MapReduce 2 and its more flexible execution model (YARN). Store large datasets with the Hadoop Distributed File System (HDFS) Run distributed computations with MapReduce Use Hadoop's data and I/O building blocks for compression, data integrity, serialization (including Avro), and persistence Discover common pitfalls and advanced features for writing real-world MapReduce programs Design, build, and administer a dedicated Hadoop cluster—or run Hadoop in the cloud Load data from relational databases into HDFS, using Sqoop Perform large-scale data processing with the Pig query language Analyze datasets with Hive, Hadoop's data

warehousing system Take advantage of HBase for structured and semi-structured data, and ZooKeeper for building distributed systems

This book presents the signal processing and data mining challenges encountered in drilling engineering, and describes the methods used to overcome them. In drilling engineering, many signal processing technologies are required to solve practical problems, such as downhole information transmission, spatial attitude of drillstring, drillstring dynamics, seismic activity while drilling, among others. This title attempts to bridge the gap between the signal processing and data mining and oil and gas drilling engineering communities. There is an urgent need to summarize signal processing and data mining issues in drilling engineering so that practitioners in these fields can understand each other in order to enhance oil and gas drilling functions. In summary, this book shows the importance of signal processing and data mining to researchers and professional drilling engineers and open up a new area of application for signal processing and data mining scientists.

Now in the 5th edition, Cracking the Coding Interview gives you the interview preparation you need to get the top software developer jobs. This book provides: 150 Programming Interview Questions and Solutions: From binary trees to binary search, this list of 150 questions includes the most common and most useful questions in data structures, algorithms, and knowledge based questions. 5 Algorithm Approaches: Stop being blind-sided by tough algorithm questions, and learn these five approaches to tackle the trickiest problems. Behind the Scenes of the interview processes at Google, Amazon, Microsoft, Facebook, Yahoo, and Apple: Learn what really goes on during your interview day and how decisions get made. Ten Mistakes Candidates Make -- And How to Avoid Them: Don't lose your dream job by making these common mistakes. Learn what many candidates do wrong, and how to avoid these issues. Steps to Prepare for Behavioral and Technical Questions: Stop meandering through an endless set of questions, while missing some of the most important preparation techniques. Follow these steps to more thoroughly prepare in less time. By examining the stories of companies like LoJack, MySpace and Groupon, an expert in Strategic Management describes how organizations can sustain high-growth through a repeatable process for innovation without succumbing to the growth hump. 30,000 first printing.

*The most updated PostgreSQL book on the market, covering version 8.0 *Highlights the most popular PostgreSQL APIs, including C, Perl, PHP, and Java *This is two books in one; it simultaneously covers key relational database design principles, while teaching PostgreSQL

The Design and Implementation of Modern Column-Oriented Database Systems discusses modern column-stores, their architecture and evolution as well the benefits they can bring in data analytics.

Unique prospective on the big data analytics phenomenon for both business and IT professionals The availability of Big Data, low-cost commodity hardware and new information management and analytics software has produced a unique moment in the history of business. The convergence of these trends means that we have the capabilities required to analyze astonishing data sets quickly and cost-effectively for the first time in history. These capabilities are neither theoretical nor trivial. They represent a genuine leap forward and a clear opportunity to realize enormous gains in terms of efficiency, productivity, revenue and profitability. The Age of Big Data is here, and these are truly revolutionary times. This timely book looks at cutting-edge companies supporting an exciting new generation of business analytics. Learn more about the trends in big data and how they are impacting the business world (Risk, Marketing, Healthcare, Financial Services, etc.) Explains this new technology and how companies can use them effectively to gather the data that they need and glean critical insights Explores relevant topics such as data privacy, data visualization, unstructured data, crowd sourcing data scientists, cloud computing for big data, and much more.

For courses in decision support systems, computerized decision-making tools, and management support systems. Market-leading guide to modern analytics, for better business decisions Analytics, Data Science, & Artificial Intelligence: Systems for Decision Support is the most comprehensive introduction to technologies collectively called analytics (or business analytics) and the fundamental methods, techniques, and software used to design and develop these systems. Students gain inspiration from examples of organisations that have employed analytics to make decisions, while leveraging the resources of a companion website. With six new chapters, the 11th edition marks a major reorganisation reflecting a new focus -- analytics and its enabling technologies, including AI, machine-learning, robotics, chatbots, and IoT.

Learn all you need to know about seven key innovations disrupting business analytics today. These innovations—the open source business model, cloud analytics, the Hadoop ecosystem, Spark and in-memory analytics, streaming analytics, Deep Learning, and self-service analytics—are radically changing how businesses use data for competitive advantage. Taken together, they are disrupting the business analytics value chain, creating new opportunities. Enterprises who seize the opportunity will thrive and prosper, while others struggle and decline: disrupt or be disrupted. Disruptive Business Analytics provides strategies to profit from disruption. It shows you how to organize for insight, build and provision an open source stack, how to practice lean data warehousing, and how to assimilate disruptive innovations into an organization. Through a short history of business analytics and a detailed survey of products and services, analytics authority Thomas W. Dinsmore provides a practical explanation of the most compelling innovations available today. What You'll Learn Discover how the open source business model works and how to make it work for you See how cloud computing completely changes the economics of analytics Harness the power of Hadoop and its

ecosystem Find out why Apache Spark is everywhere Discover the potential of streaming and real-time analytics Learn what Deep Learning can do and why it matters See how self-service analytics can change the way organizations do business Who This Book Is For Corporate actors at all levels of responsibility for analytics: analysts, CIOs, CTOs, strategic decision makers, managers, systems architects, technical marketers, product developers, IT personnel, and consultants.

Joe Celko's SQL Puzzles and Answers Elsevier

In this insightful book, you'll learn from the best data practitioners in the field just how wide-ranging -- and beautiful -- working with data can be. Join 39 contributors as they explain how they developed simple and elegant solutions on projects ranging from the Mars lander to a Radiohead video. With Beautiful Data, you will: Explore the opportunities and challenges involved in working with the vast number of datasets made available by the Web Learn how to visualize trends in urban crime, using maps and data mashups Discover the challenges of designing a data processing system that works within the constraints of space travel Learn how crowdsourcing and transparency have combined to advance the state of drug research Understand how new data can automatically trigger alerts when it matches or overlaps pre-existing data Learn about the massive infrastructure required to create, capture, and process DNA data That's only small sample of what you'll find in Beautiful Data. For anyone who handles data, this is a truly fascinating book. Contributors include: Nathan Yau Jonathan Follett and Matt Holm J.M. Hughes Raghu Ramakrishnan, Brian Cooper, and Utkarsh Srivastava Jeff Hammerbacher Jason Dykes and Jo Wood Jeff Jonas and Lisa Sokol Jud Valeski Alon Halevy and Jayant Madhavan Aaron Koblin with Valdean Klump Michal Migurski Jeff Heer Coco Krumme Peter Norvig Matt Wood and Ben Blackburne Jean-Claude Bradley, Rajarshi Guha, Andrew Lang, Pierre Lindenbaum, Cameron Neylon, Antony Williams, and Egon Willighagen Lukas Biewald and Brendan O'Connor Hadley Wickham, Deborah Swayne, and David Poole Andrew Gelman, Jonathan P. Kestellec, and Yair Ghitza Toby Segaran

The debut cookbook from the popular New York Times website and mobile app NYT Cooking, featuring 100 vividly photographed no-recipe recipes to make weeknight cooking more inspired and delicious. You don't need a recipe. Really, you don't. Sam Sifton, founding editor of New York Times Cooking, makes improvisational cooking easier than you think. In this handy book of ideas, Sifton delivers more than one hundred no-recipe recipes—each gloriously photographed—to make with the ingredients you have on hand or could pick up on a quick trip to the store. You'll see how to make these meals as big or as small as you like, substituting ingredients as you go. Fried Egg Quesadillas. Pizza without a Crust. Weeknight Fried Rice. Pasta with Garbanzos. Roasted Shrimp Tacos. Chicken with Caramelized Onions and Croutons. Oven S'Mores. Welcome home to freestyle, relaxed cooking that is absolutely yours.

Data Warehousing in the Age of the Big Data will help you and your organization make the most of unstructured data with your existing data warehouse. As Big Data continues to revolutionize how we use data, it doesn't have to create more confusion. Expert author Krish Krishnan helps you make sense of how Big Data fits into the world of data warehousing in clear and concise detail.

The book is presented in three distinct parts. Part 1 discusses Big Data, its technologies and use cases from early adopters. Part 2

addresses data warehousing, its shortcomings, and new architecture options, workloads, and integration techniques for Big Data and the data warehouse. Part 3 deals with data governance, data visualization, information life-cycle management, data scientists, and implementing a Big Data-ready data warehouse. Extensive appendixes include case studies from vendor implementations and a special segment on how we can build a healthcare information factory. Ultimately, this book will help you navigate through the complex layers of Big Data and data warehousing while providing you information on how to effectively think about using all these technologies and the architectures to design the next-generation data warehouse. Learn how to leverage Big Data by effectively integrating it into your data warehouse. Includes real-world examples and use cases that clearly demonstrate Hadoop, NoSQL, HBASE, Hive, and other Big Data technologies Understand how to optimize and tune your current data warehouse infrastructure and integrate newer infrastructure matching data processing workloads and requirements

This book constitutes selected papers from the 14th European, Mediterranean, and Middle Eastern Conference, EMCIS 2017, held in Coimbra, Portugal, in September 2017. EMCIS is focusing on approaches that facilitate the identification of innovative research of significant relevance to the IS discipline following sound research methodologies that lead to results of measurable impact. The 37 full and 16 short papers presented in this volume were carefully reviewed and selected from a total of 106 submissions. They are organized in sections on big data and Semantic Web; digital services, social media and digital collaboration; e-government; healthcare information systems; information systems security and information privacy protection; IT governance; and management and organizational issues in information systems.

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

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!

PostgreSQL offers a comprehensive set of replication related features. Unleashing the power of PostgreSQL provides you with countless opportunities and a competitive advantage over other database systems. This book will guide you through the most important concepts of PostgreSQL replication. It contains all the information you need to design and operate replicated setups. Beginning by giving you an understanding of replication concepts, the PostgreSQL transaction log, and Point-in-time Recovery, we gradually move on to setting up asynchronous and synchronous replication. Next up, you will learn to monitor a PostgreSQL cluster setup, deal with monitoring tools, and then move on to understanding Linux High Availability. Further, we explore widely-used tools such as Slony, SkyTools, Postgres-XC, and walbouncer, and set up PL/Proxy. Finally, you'll get acquainted with the new technology of BDR, which allows bidirectional replication in PostgreSQL.

The child of a commoner family in Su Liuxia. After graduating from university, she didn't study in her own department, but instead went to her favorite performing arts field. Because he grew up in the countryside, his family wanted him to be a teacher or doctor, but he didn't like it. In the end, he studied economics and management. Li Youjin, a CEO of a listed company, was cold, quiet and dark. After he left the country to study abroad, he set up the current J-Y group with his friends. "At the age of 25, he did not have a single woman by his side. The people outside also said that he liked men, and the family was afraid of that, so they arranged a blind date for him. Since he did not want to accept the arrangements, he could only flee." It's fine. I'll raise you in the future. At worst, you just have to take on more tricks! "Su Liuxia patted her chest." ... " "Good," Liao Jin smiled as he looked at her righteous demeanor. Is there still time for me to go back on my word? " Su Liuxia, who had already arrived at the Li residence, looked at Li Youjin and spoke again. Hearing ... Join Collection

Know how your company can accelerate growth by not only tapping into new growth vectors, but also by adapting its organization, culture, and processes. To oversee growth from an idea to a company with billions in revenue, CEOs must reinvent many aspects of their company in anticipation of it reaching ever-higher revenues. Author Peter Cohan takes you through the four stages of scaling: winning the first customers, building a scalable business model, sprinting to liquidity, and running the marathon. What You'll Learn Discover how founders keep their CEO positions by managing the organizational change needed to reach the next stage of scaling Read case studies that illustrate how CEOs craft growth strategies, raise capital, create culture, build their

organizations, set goals, and manage processes to achieve them Discover principles of successful scaling through comparisons of successful and less successful companies Use the Scaling Quotient to assess your startup's readiness to grow Follow a road map for turning your idea into a company that can change the world Who This Book Is For Entrepreneurs, aspiring CEOs, capital providers, and all other key stakeholders

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. Mastering PostgreSQL in Application Development is intended for developers working on applications that use a database server. The book addresses specifically the PostgreSQL RDBMS: it actually is the world's most advanced Open Source database as said in its slogan on the official website. By the end of this book, you will know why, and agree!

This collection represents the full spectrum of data-related content we've published on O'Reilly Radar over the last year. Mike Loukides kicked things off in June 2010 with "What is data science?" and from there we've pursued the various threads and themes that naturally emerged. Now, roughly a year later, we can look back over all we've covered and identify a number of core data areas: Data issues -- The opportunities and ambiguities of the data space are evident in discussions around privacy, the implications of data-centric industries, and the debate about the phrase "data science" itself. The application of data: products and processes -- A "data product" can emerge from virtually any domain, including everything from data startups to established enterprises to media/journalism to education and research. Data science and data tools -- The tools and technologies that drive data science are of course essential to this space, but the varied techniques being applied are also key to understanding the big data arena. The business of data -- Take a closer look at the actions connected to data -- the finding, organizing, and analyzing that provide organizations of all sizes with the information they need to compete.

Mastering Cloud Computing is designed for undergraduate students learning to develop cloud computing applications. Tomorrow's applications won't live on a single computer but will be deployed from and reside on a virtual server, accessible anywhere, any

time. Tomorrow's application developers need to understand the requirements of building apps for these virtual systems, including concurrent programming, high-performance computing, and data-intensive systems. The book introduces the principles of distributed and parallel computing underlying cloud architectures and specifically focuses on virtualization, thread programming, task programming, and map-reduce programming. There are examples demonstrating all of these and more, with exercises and labs throughout. Explains how to make design choices and tradeoffs to consider when building applications to run in a virtual cloud environment Real-world case studies include scientific, business, and energy-efficiency considerations

This book celebrates Michael Stonebraker's accomplishments that led to his 2014 ACM A.M. Turing Award "for fundamental contributions to the concepts and practices underlying modern database systems." The book describes, for the broad computing community, the unique nature, significance, and impact of Mike's achievements in advancing modern database systems over more than forty years. Today, data is considered the world's most valuable resource, whether it is in the tens of millions of databases used to manage the world's businesses and governments, in the billions of databases in our smartphones and watches, or residing elsewhere, as yet unmanaged, awaiting the elusive next generation of database systems. Every one of the millions or billions of databases includes features that are celebrated by the 2014 Turing Award and are described in this book. Why should I care about databases? What is a database? What is data management? What is a database management system (DBMS)? These are just some of the questions that this book answers, in describing the development of data management through the achievements of Mike Stonebraker and his over 200 collaborators. In reading the stories in this book, you will discover core data management concepts that were developed over the two greatest eras (so far) of data management technology. The book is a collection of 36 stories written by Mike and 38 of his collaborators: 23 world-leading database researchers, 11 world-class systems engineers, and 4 business partners. If you are an aspiring researcher, engineer, or entrepreneur you might read these stories to find these turning points as practice to tilt at your own computer-science windmills, to spur yourself to your next step of innovation and achievement. Learn Rails the way the Rails core team recommends it, along with the tens of thousands of developers who have used this broad, far-reaching tutorial and reference. If you're new to Rails, you'll get step-by-step guidance. If you're an experienced developer, get the comprehensive, insider information you need for the latest version of Ruby on Rails. The new edition of this award-winning classic is completely updated for Rails 6 and Ruby 2.6, with information on system testing, Webpack, and advanced JavaScript. Ruby on Rails helps you produce high-quality, beautiful-looking web applications quickly - you concentrate on creating the application, and Rails takes care of the details. Rails 6 brings many improvements, and this edition is updated to cover the new features and changes in best practices. We start with a step-by-step walkthrough of building a real application, and in-depth chapters look at the built-in Rails features. Follow along with an extended tutorial as you write a web-based store application. Eliminate tedious configuration and housekeeping, seamlessly incorporate Ajax and JavaScript, send and receive emails, manage background jobs with ActiveJob, and

build real-time features using WebSockets and ActionCable. Test your applications as you write them using the built-in unit, integration, and system testing frameworks, internationalize your applications, and deploy your applications easily and securely. New in this edition is coverage of Action Mailer, which allows you to receive emails in your app as well as ActionText, a zero-configuration rich text editing feature. Rails 1.0 was released in December 2005. This book was there from the start, and didn't just evolve alongside Rails, it evolved with Rails. It has been developed in consultation with the Rails core team. In fact, Rails itself is tested against the code in this book. What You Need: All you need is a Windows, Mac OS X, or Linux machine to do development on. This book will take you through the steps to install Rails and its dependencies. If you aren't familiar with the Ruby programming language, this book contains a chapter that covers the basics necessary to understand the material in the book.

In Romania at the height of Ceausescu's reign, several young people leave the impoverished provinces for the city in search of better prospects, but they must face betrayal, suicide, and the reality that even the strongest must bend to the oppressors or resist and die.

This book presents and discusses the main strategic and organizational challenges posed by Big Data and analytics in a manner relevant to both practitioners and scholars. The first part of the book analyzes strategic issues relating to the growing relevance of Big Data and analytics for competitive advantage, which is also attributable to empowerment of activities such as consumer profiling, market segmentation, and development of new products or services. Detailed consideration is also given to the strategic impact of Big Data and analytics on innovation in domains such as government and education and to Big Data-driven business models. The second part of the book addresses the impact of Big Data and analytics on management and organizations, focusing on challenges for governance, evaluation, and change management, while the concluding part reviews real examples of Big Data and analytics innovation at the global level. The text is supported by informative illustrations and case studies, so that practitioners can use the book as a toolbox to improve understanding and exploit business opportunities related to Big Data and analytics.

The Campaign is a close-up look at the paranoid, frenzied, oppressive, and exhilarating world of modern political campaigns? a universe where truth is fungible and moral conviction a mere asset, like good looks or personal wealth. Corporeal restraints do not exist. People regularly become things they are not. Evan Mandery, research director on Ruth Messinger's doomed challenge to Mayor Rudy Giuliani, offers a behind-the-scenes look at political campaigns in the television era. A day-to-day account of the 1997 New York City mayoral race, it takes us to the real battlegrounds of modern politics: polls, focus groups and television editing studios. With Mandery as our guide, we watch first-hand as political consultants, conceive of the ideal candidate and then attempt to fit their client into that ideal, no matter how

uncomfortably. The stars of the story are memorable: Rudy Giuliani, popping his eyes and tweaking the truth; Al Sharpton, the colorful preacher and rising political force; and Ruth Messinger herself, torn between her populist political upbringing and the modern political world where money dominates over all other concerns. Sometimes cynical, often mirthful, and always honest, *The Campaign* will forever change your view of political campaigns.

You Let Me In delivers a stunning tale from debut author Camilla Bruce, combining the sinister domestic atmosphere of Gillian Flynn's *Sharp Objects* with the otherworldly thrills of Neil Gaiman's *The Ocean at the End of the Lane*. Cassandra Tipp is dead...or is she? After all, the notorious recluse and eccentric bestselling novelist has always been prone to flights of fancy—everyone in town remembers the shocking events leading up to Cassie's infamous trial (she may have been acquitted, but the insanity defense only stretches so far). Cassandra Tipp has left behind no body—just her massive fortune, and one final manuscript. Then again, there are enough bodies in her past—her husband Tommy Tipp, whose mysterious disembowelment has never been solved, and a few years later, the shocking murder-suicide of her father and brother. Cassandra Tipp will tell you a story—but it will come with a terrible price. What really happened, out there in the woods—and who has Cassie been protecting all along? Read on, if you dare... At the Publisher's request, this title is being sold without Digital Rights Management Software (DRM) applied.

This book contains a number of chapters on transactional database concurrency control. A two-sentence summary of the volume's entire sequence of chapters is this: traditional locking techniques can be improved in multiple dimensions, notably in lock scopes (sizes), lock modes (increment, decrement, and more), lock durations (late acquisition, early release), and lock acquisition sequence (to avoid deadlocks). Even if some of these improvements can be transferred to optimistic concurrency control, notably a fine granularity of concurrency control with serializable transaction isolation including phantom protection, pessimistic concurrency control is categorically superior to optimistic concurrency control, i.e., independent of application, workload, deployment, hardware, and software implementation.

"Regardless of the type of architecture, architecture itself is an organized accumulation of knowledge within a particular domain. While we generally conceive of its representation as a set of diagrams, containing specific notations and taxonomies of symbols and glossary terms, an architecture may actually be represented using anything that can be arranged in a pattern to record information. The earliest forms of architecture relate to architecting buildings, monuments, military disciplines, organized religion, music, storytelling, and various other forms within the arts. These early forms of architecture of course predate computer related architectures by thousands of years. That said, it is worth noting that there are a number of common elements among architectures irrespective of their relative age, such as forms of standardization, reusable structures, the accumulation of knowledge, and providing a context for understanding

something. Needless to say, anyone can be an architect in a topic in which they have a deep understanding and appreciation of. While one obvious difference among architects is the amount and variety of pertinent experience, the less obvious difference is the degree to which an architect recognizes the potential forms of standardization, reusable structures, accumulation of knowledge, relationships among the components, and use of architecture as an accelerator to more rapidly understand the context and scope of a particular topic or to rapidly convey it to another. Architectures as a result must be easy to understand"--

Learn what it takes to succeed in the the most in-demand tech job Harvard Business Review calls it the sexiest tech job of the 21st century. Data scientists are in demand, and this unique book shows you exactly what employers want and the skill set that separates the quality data scientist from other talented IT professionals. Data science involves extracting, creating, and processing data to turn it into business value. With over 15 years of big data, predictive modeling, and business analytics experience, author Vincent Granville is no stranger to data science. In this one-of-a-kind guide, he provides insight into the essential data science skills, such as statistics and visualization techniques, and covers everything from analytical recipes and data science tricks to common job interview questions, sample resumes, and source code. The applications are endless and varied: automatically detecting spam and plagiarism, optimizing bid prices in keyword advertising, identifying new molecules to fight cancer, assessing the risk of meteorite impact. Complete with case studies, this book is a must, whether you're looking to become a data scientist or to hire one. Explains the finer points of data science, the required skills, and how to acquire them, including analytical recipes, standard rules, source code, and a dictionary of terms Shows what companies are looking for and how the growing importance of big data has increased the demand for data scientists Features job interview questions, sample resumes, salary surveys, and examples of job ads Case studies explore how data science is used on Wall Street, in botnet detection, for online advertising, and in many other business-critical situations Developing Analytic Talent: Becoming a Data Scientist is essential reading for those aspiring to this hot career choice and for employers seeking the best candidates.

[Copyright: 2fd40e68ffb972b8dba05da954524148](https://www.dwgeek.com/2fd40e68ffb972b8dba05da954524148)