

Database System Using Oracle Nilesh Shah

Probabilistic databases are databases where the value of some attributes or the presence of some records are uncertain and known only with some probability. Applications in many areas such as information extraction, RFID and scientific data management, data cleaning, data integration, and financial risk assessment produce large volumes of uncertain data, which are best modeled and processed by a probabilistic database. This book presents the state of the art in representation formalisms and query processing techniques for probabilistic data. It starts by discussing the basic principles for representing large probabilistic databases, by decomposing them into tuple-independent tables, block-independent-disjoint tables, or U-databases. Then it discusses two classes of techniques for query evaluation on probabilistic databases. In extensional query evaluation, the entire probabilistic inference can be pushed into the database engine and, therefore, processed as effectively as the evaluation of standard SQL queries. The relational queries that can be evaluated this way are called safe queries. In intensional query evaluation, the probabilistic inference is performed over a propositional formula called lineage expression: every relational query can be evaluated this way, but the data complexity dramatically depends on the query being evaluated, and can be #P-hard. The book also discusses some advanced topics in probabilistic data management such as top-k query processing, sequential probabilistic databases, indexing and materialized views, and Monte Carlo databases. Table of Contents: Overview / Data and Query Model / The Query Evaluation Problem / Extensional Query Evaluation / Intensional Query Evaluation / Advanced Techniques

Updated to include the new features introduced in PostgreSQL 13, this book shows you how to build better PostgreSQL applications and administer your PostgreSQL database efficiently. You'll master the advanced features of PostgreSQL and develop the skills you need to build secure and highly available database solutions.

Learn the concepts, principles, design, implementation, and management issues of databases. You will adopt a methodical and pragmatic approach to solving database systems problems. Database Systems: A Pragmatic Approach provides a comprehensive, yet concise introduction to database systems, with special emphasis on the relational database model. This book discusses the database as an essential component of a software system, as well as a valuable, mission-critical corporate resource. New in this second edition is updated SQL content covering the latest release of the Oracle Database Management System along with a reorganized sequence of the topics which is more useful for learning. Also included are revised and additional illustrations, as well as a new chapter on using relational databases to anchor large, complex management support systems. There is also added reference content in the appendixes. This book is based on lecture notes that have been tested and proven over several years, with outstanding results. It combines a balance of theory with practice, to give you your best chance at success. Each chapter is organized systematically into brief sections, with itemization of the important points to be remembered. Additionally, the book includes a number of author Elvis Foster's original methodologies that add clarity and creativity to the database modeling and design experience. What You'll Learn Understand the relational model and the advantages it brings to software systems Design database schemas with integrity rules that ensure correctness of corporate data Query data using SQL in order to generate reports, charts, graphs, and other business results Understand what it means to be a database administrator, and why the profession is highly paid Build and manage web-accessible databases in support of applications delivered via a browser Become familiar with the common database brands, their similarities and differences Explore special topics such as tree-based data, hashing for fast access, distributed and object databases, and more Who This Book Is For Students who are studying database technology, who aspire to a career as a database administrator or designer, and practicing database administrators and developers desiring to strengthen their knowledge of database theory

Master one of the world's most powerful enterprise workload automation tools? BMC Control-M 7 - using this book and eBook. This manual is specially written for Students who are interested in understanding Structured Query Language and PL-SQL concepts in the Computer Engineering and Information technology field and wants to gain enhance knowledge about power of SQL Language in Relational Database Management System Development. The manual covers practical point of view in all aspects of SQL and PL/SQL including DDL, DML, DCL sublanguages, also there are practices for Views, Group by, Having Clause. All PL-SQL concepts like Condition and Loop Structures, Functions and Procedures, Cursor, Triggers, Locks are illustrated using best examples

The book shows you how to view data from multiple perspectives, including data frame and column attributes. You will cover common and not-so-common challenges that are faced while cleaning messy data for complex situations. You will learn to manipulate data and get them down to a form that can be useful for making the right decisions.

This advanced graduate textbook gives an authoritative and insightful description of the major ideas and techniques of public key cryptography.

Find out how to craft effective, business-oriented Java EE 8 applications that target customer's demands in the age of Cloud platforms and container technology. About This Book Understand the principles of modern Java EE and how to realize effective architectures Gain knowledge of how to design enterprise software in the age of automation, Continuous Delivery and Cloud platforms Learn about the reasoning and motivations behind state-of-the-art enterprise Java technology, that focuses on business Who This Book Is For This book is for experienced Java EE developers who are aspiring to become the architects of enterprise-grade applications, or software architects who would like to leverage Java EE to create effective blueprints of applications. What You Will Learn What enterprise software engineers should focus on Implement applications, packages, and components in a modern way Design and structure application architectures Discover how to realize technical and cross-cutting aspects Get to grips with containers and container orchestration technology Realize zero-dependency, 12-factor, and Cloud-native applications Implement automated, fast, reliable, and maintainable software tests Discover distributed system architectures and their requirements In Detail Java EE 8 brings with it a load of features, mainly targeting newer architectures such as microservices, modernized security APIs, and cloud deployments. This book will teach you to design and develop modern, business-oriented applications using Java EE 8. It shows how to structure systems and applications, and how design patterns and Domain Driven Design aspects are realized in the age of Java EE 8. You will learn about the concepts and principles behind Java EE applications, and how to effect communication, persistence, technical and cross-cutting concerns, and asynchronous behavior. This book covers Continuous Delivery, DevOps, infrastructure-as-code, containers, container orchestration technologies, such as Docker

and Kubernetes, and why and especially how Java EE fits into this world. It also covers the requirements behind containerized, zero-dependency applications and how modern Java EE application servers support these approaches. You will also learn about automated, fast, and reliable software tests, in different test levels, scopes, and test technologies. This book covers the prerequisites and challenges of distributed systems that lead to microservice, shared-nothing architectures. The challenges and solutions of consistency versus scalability will further lead us to event sourcing, event-driven architectures, and the CQRS principle. This book also includes the nuts and bolts of application performance as well as how to realize resilience, logging, monitoring and tracing in a modern enterprise world. Last but not least the demands of securing enterprise systems are covered. By the end, you will understand the ins and outs of Java EE so that you can make critical design decisions that not only live up to, but also surpass your clients' expectations. Style and approach This book focuses on solving business problems and meeting customer demands in the enterprise world. It covers how to create enterprise applications with reasonable technology choices, free of cargo-cult and over-engineering. The aspects shown in this book not only demonstrate how to realize a certain solution, but also explain its motivations and reasoning.

Written as a tutorial to explore and understand the power of R for machine learning. This practical guide that covers all of the need to know topics in a very systematic way. For each machine learning approach, each step in the process is detailed, from preparing the data for analysis to evaluating the results. These steps will build the knowledge you need to apply them to your own data science tasks. Intended for those who want to learn how to use R's machine learning capabilities and gain insight from your data. Perhaps you already know a bit about machine learning, but have never used R; or perhaps you know a little R but are new to machine learning. In either case, this book will get you up and running quickly. It would be helpful to have a bit of familiarity with basic programming concepts, but no prior experience is required.

Introduction -- China's Sputnik moment -- Copycats in the Coliseum -- China's alternate Internet universe -- A tale of two countries -- The four waves of AI -- Utopia, dystopia, and the real AI crisis -- The wisdom of cancer -- A blueprint for human co-existence with AI -- Our global AI story

No IT server platform is 100% secure and useful at the same time. If your server is installed in a secure vault, three floors underground in a double-locked room, not connected to any network and switched off, one would say it was reasonably secure, but it would be a stretch to call it useful. This IBM® Redbooks® publication is about switching on the power to your Linux® on System z® server, connecting it to the data and to the network, and letting users have access to this formidable resource space in a secure, controlled, and auditable fashion to make sure the System z server and Linux are useful to your business. As the quotation illustrates, the book is also about ensuring that, before you start designing a security solution, you understand what the solution has to achieve. The base for a secure system is tightly related to the way the architecture and virtualization has been implemented on IBM System z. Since its inception 45 years ago, the architecture has been continuously developed to meet the increasing demands for a more secure and stable platform. This book is intended for system engineers and security administrators who want to customize a Linux on System z environment to meet strict security, audit, and control regulations. For additional information, there is a tech note that describes the best practices for securing your network. It can be found at:

<http://www.redbooks.ibm.com/abstracts/tips0981.html?Open>

Query Processing on Probabilistic Data: A Survey presents the main approaches developed in the literature on probabilistic relational data, reconciling concepts developed in parallel by the Database and Artificial Intelligence communities.

Create and unleash the power of neural networks by implementing professional, clean, and clear Java code About This Book* Learn to build amazing projects using neural networks including forecasting the weather and pattern recognition* Explore the Java multi-platform feature to run your personal neural networks everywhere* This step-by-step guide will help you solve real-world problems and links neural network theory to their application Who This Book Is For This book is for Java developers who want to know how to develop smarter applications using the power of neural networks. Those who deal with a lot of complex data and want to use it efficiently in their day-to-day apps will find this book quite useful. Some basic experience with statistical computations is expected. What You Will Learn* Develop an understanding of neural networks and how they can be fitted* Explore the learning process of neural networks* Build neural network applications with Java using hands-on examples* Discover the power of neural network's unsupervised learning process to extract the intrinsic knowledge hidden behind the data* Apply the code generated in practical examples, including weather forecasting and pattern recognition* Understand how to make the best choice of learning parameters to ensure you have a more effective application* Select and split data sets into training, test, and validation, and explore validation strategies In Detail Want to discover the current state-of-art in the field of neural networks that will let you understand and design new strategies to apply to more complex problems? This book takes you on a complete walkthrough of the process of developing basic to advanced practical examples based on neural networks with Java, giving you everything you need to stand out. You will first learn the basics of neural networks and their process of learning. We then focus on what Perceptrons are and their features. Next, you will implement self-organizing maps using practical examples. Further on, you will learn about some of the applications that are presented in this book such as weather forecasting, disease diagnosis, customer profiling, generalization, extreme machine learning, and characters recognition (OCR). Finally, you will learn methods to optimize and adapt neural networks in real time. All the examples generated in the book are provided in the form of illustrative source code, which merges object-oriented programming (OOP) concepts and neural network features to enhance your learning experience.

Design RPA solutions to perform a wide range of transactional tasks with minimal cost and maximum ROI Key Features A beginner's guide to learn Robotic Process Automation and its impact on the modern world Design, test, and perform enterprise automation task with UiPath Create Automation apps and deploy them to all the computers in your

department. Book Description Robotic Process Automation (RPA) enables automating business processes using software robots. Software robots interpret, trigger responses, and communicate with other systems just like humans do. Robotic processes and intelligent automation tools can help businesses improve the effectiveness of services faster and at a lower cost than current methods. This book is the perfect start to your automation journey, with a special focus on one of the most popular RPA tools: UiPath. Learning Robotic Process Automation takes you on a journey from understanding the basics of RPA to advanced implementation techniques. You will become oriented in the UiPath interface and learn about its workflow. Once you are familiar with the environment, we will get hands-on with automating different applications such as Excel, SAP, Windows and web applications, screen and web scraping, working with user events, as well as understanding exceptions and debugging. By the end of the book, you'll not only be able to build your first software bot, but also you'll wire it to perform various automation tasks with the help of best practices for bot deployment. What you will learn Understand Robotic Process Automation technology Learn UiPath programming techniques to deploy robot configurations Explore various data extraction techniques Learn about integrations with various popular applications such as SAP and MS Office Debug a programmed robot including logging and exception handling Maintain code version and source control Deploy and control Bots with UiPath Orchestrator Who this book is for If you would like to pursue a career in Robotic Process Automation or improve the efficiency of your businesses by automating common tasks, then this book is perfect for you. Prior programming knowledge of either Visual Basic or C# will be useful.

Presents the fundamental concepts of database management. This text is suitable for a first course in databases at the junior/senior undergraduate level or the first year graduate level.

This text provides comprehensive instruction that should enable students to develop a database in Oracle, using the database and application development tools provided by the Oracle Academic Alliance Program. It also provides instructors and technical support personnel with instruction on administration of these tools.

See how Domain-Driven Design (DDD) combines with Jakarta EE MicroProfile or Spring Boot to offer a complete suite for building enterprise-grade applications. In this book you will see how these all come together in one of the most efficient ways to develop complex software. Practical Domain-Driven Design in Enterprise Java starts by building out the Cargo Tracker reference application as a monolithic application using the Jakarta EE platform. By doing so, you will map concepts of DDD (bounded contexts, language, and aggregates) to the corresponding available tools (CDI, JAX-RS, and JPA) within the Jakarta EE platform. Once you have completed the monolithic application, you will walk through the complete conversion of the monolith to a microservices-based architecture, again mapping the concepts of DDD and the corresponding available tools within the MicroProfile platform (config, discovery, and fault tolerance). To finish this section, you will examine the same microservices architecture on the Spring Boot platform. The final set of chapters looks at what the application would be like if you used the CQRS and event sourcing patterns. Here you'll use the Axon framework as the base framework. What You Will Learn Discover the DDD architectural principles and use the DDD design patterns Use the new Eclipse Jakarta EE platform Work with the Spring Boot framework Implement microservices design patterns, including context mapping, logic design, entities, integration, testing, and security Carry out event sourcing Apply CQRS Who This Book Is For Junior developers intending to start working on enterprise Java; senior developers transitioning from monolithic- to microservices-based architectures; and architects transitioning to a DDD philosophy of building applications.

Provides information on ways to break into and defend seven database servers, covering such topics as identifying vulnerabilities, how an attack is carried out, and how to stop an attack.

Learn to design, build, and manage your infrastructure on the most popular of all the Cloud platforms—Amazon Web Services About This Book Learn how to leverage various Amazon Web Services (AWS) components and services to build a secure, reliable, and robust environment to host your applications on Deep dive into the core AWS service offerings with hands-on tutorials, real-world use case scenarios, and best practices A self-paced, systematic, and step-by-step guide to learning and implementing AWS in your own environment Who This Book Is For This book is for those who want to learn and leverage AWS. Although no prior experience with AWS is required, it is recommended that you have some hands-on experience of Linux, Web Services, and basic networking What You Will Learn A brief introduction to Cloud Computing and AWS accompanied by steps to sign up for your first AWS account Create and manage users, groups, and permissions using AWS Identity and Access Management services Get started with deploying and accessing EC2 instances, working with EBS Volumes and Snapshots Customize and create your very own Amazon Machine Image Design and deploy your instances on a highly secured, network isolated environment using Amazon VPC Effectively monitor your AWS environment using specialized alarms, custom monitoring metrics, and much more Explore the various benefits of Database-as-a-Service offerings and leverage them using Amazon RDS and Amazon DynamoDB Take an in-depth look at what's new with AWS, including EC2 Container Service and Elastic File System In Detail AWS is at the forefront of Cloud Computing today. Many businesses are moving away from traditional datacenters and toward AWS because of its reliability, vast service offerings, lower costs, and high rate of innovation. Because of its versatility and flexible design, AWS can be used to accomplish a variety of simple and complicated tasks such as hosting multitier websites, running large scale parallel processing, content delivery, petabyte storage and archival, and lots more. Whether you are a seasoned sysadmin or a rookie, this book will provide you with all the necessary skills to design, deploy, and manage your applications on the AWS cloud platform. The book guides you through the core AWS services such as IAM, EC2, VPC, RDS, and S3 using a simple real world application hosting example that you can relate to. Each chapter is designed to provide you with the most information possible about a particular AWS service coupled with easy to follow hands-on steps, best practices, tips, and recommendations. By the end of the book, you will be able to create a highly secure, fault tolerant, and scalable environment for your applications to run on. Style and approach This in-depth and insightful guide is filled with easy-to-follow examples, real-world use cases, best practices, and recommendations that will help you design and leverage AWS.

Database Systems: A Pragmatic Approach is a classroom textbook for use by students who are learning about relational databases, and the professors who teach them. It discusses the database as an essential component of a software system, as well as a valuable, mission critical corporate resource. The book is based on lecture notes that have been tested and proven over several years, with outstanding results. It also exemplifies mastery of the technique of combining and balancing theory with practice, to give students their best chance at success.

Upholding his aim for brevity, comprehensive coverage, and relevance, author Elvis C. Foster's practical and methodical discussion style gets straight to the salient issues, and avoids unnecessary fluff as well as an overkill of theoretical calculations. The book discusses concepts, principles, design, implementation, and management issues of databases. Each chapter is organized systematically into brief, reader-friendly sections, with itemization of the important points to be remembered. It adopts a methodical and pragmatic approach to solving database systems problems. Diagrams and illustrations also sum up the salient points to enhance learning. Additionally, the book includes a number of Foster's original methodologies that add clarity and creativity to the database modeling and design experience while making a novel

contribution to the discipline. Everything combines to make Database Systems: A Pragmatic Approach an excellent textbook for students, and an excellent resource on theory for the practitioner.

This book will help you get you up to speed with developing performant software robots to automate business processes. With the help of examples, practical case studies, and step-by-step instructions, you'll be able to make the most of Automation Anywhere for building RPA solutions.

The volume contains 75 papers presented at International Conference on Communication and Networks (COMNET 2015) held during February 19–20, 2016 at Ahmedabad Management Association (AMA), Ahmedabad, India and organized by Computer Society of India (CSI), Ahmedabad Chapter, Division IV and Association of Computing Machinery (ACM), Ahmedabad Chapter. The book aims to provide a forum to researchers to propose theory and technology on the networks and services, share their experience in IT and telecommunications industries and to discuss future management solutions for communication systems, networks and services. It comprises of original contributions from researchers describing their original, unpublished, research contribution. The papers are mainly from 4 areas – Security, Management and Control, Protocol and Deployment, and Applications. The topics covered in the book are newly emerging algorithms, communication systems, network standards, services, and applications.

This easy-to-read book provides quick lessons on relational database terminology and normalization with very little effort. Updated for Oracle 9i, its thorough coverage of Oracle's SQL and PL/SQL and introduction to advanced SQL topics makes this a must for busy professionals.

The many examples, with output shown as screenshots, provide ample opportunity for the reader to easily understand and learn to use Oracle and SQL. KEY TOPICS: First introducing relational database concepts, the book covers SQL (Structured Query Language); Programming Language (the extension to SQL); and then proceeds to advanced topics, which include Oracle architecture and database administration with enterprise tools. MARKET: For any IT professional who needs to understand SQL or Oracle database systems.

This IBM® Redbooks® publication examines the IBM Tivoli® Directory Server for z/OS®. IBM Tivoli Directory Server is a powerful Lightweight Directory Access Protocol (LDAP) infrastructure that provides a foundation for deploying comprehensive identity management applications and advanced software architectures. This publication provides an introduction to the IBM Tivoli Directory Server for z/OS that provides a brief summary of its features and a examination of the possible deployment topologies. It discusses planning a deployment of IBM Tivoli Directory Server for z/OS, which includes prerequisites, planning considerations, and data stores, and provides a brief overview of the configuration process. Additional chapters provide a detailed discussion of the IBM Tivoli Directory Server for z/OS architecture that examines the supported back ends, discusses in what scenarios they are best used, and provides usage examples for each back end. The discussion of schemas breaks down the schema and provides guidance on extending it. A broad discussion of authentication, authorization, and security examines the various access protections, bind mechanisms, and transport security available with IBM Tivoli Directory Server for z/OS. This chapter also provides an examination of the new Password Policy feature. Basic and advanced replication topologies are also covered. A discussion on plug-ins provides details on the various types of plug-ins, the plug-in architecture, and creating a plug-in, and provides an example plug-in. Integration of IBM Tivoli Directory Server for z/OS into the IBM Workload Manager environment is also covered. This publication also provides detailed information about the configuration of IBM Tivoli Directory Server for z/OS. It discusses deploying IBM Tivoli Directory Server for z/OS on a single system, with examples of configuring the available back ends. Configuration examples are also provided for deploying the server in a Sysplex, and for both basic and advanced replication topologies. Finally it provides guidance on monitoring and debugging IBM Tivoli Directory Server for z/OS.

Become a full-fledged Qlik Sense Consultant with the help of this unique guide About This Book Become a successful Qlik consultant with the help of this insightful guide Build what is in line as well as exceeding your customer's expectations from your Qlik Sense solutions using this highly practical guide Build result-driven optimized BI solutions using Qlik with the help of industry examples Who This Book Is For If you have basic familiarity with Qlik Sense and want to upgrade your skills to become a full-fledged Qlik Consultant, this book is for you. With this book, you will be able to create efficient business intelligence solutions that would fetch client satisfaction, and in turn, more projects. What You Will Learn Understand the importance and expectations of a consultant's role Engage with the customer to understand the ir goals and future objectives Design the optimum architecture, using the best practices for the development and implementation of your projects Ensure successful adoption using real-life examples to make your learning complete Learn about the important stages of a Qlik project's life cycle In Detail Qlik Sense is a leading platform for business intelligence (BI) solutions. Qlik Sense helps organizations in making informed decisions based on the data they have. This book will teach you how to effectively use Qlik for optimum customer satisfaction. You will undergo a metamorphosis from a developer to a consultant who is capable of building the most suitable BI solutions for your clients. The book will take you through several business cases – this will give you enough insight to understand the needs of the client clearly and build a BI solution that meets or exceeds their expectations. Starting from the pre-project activities, you will go to the actual execution of the project, the implementation, and even maintenance. This book will give you all the information you need - from the strategy to requirement gathering to implementing BI solutions using Qlik Sense. The book will empower you to take the right decisions in tricky and diffi cult situations while developing analytics and dashboards. Style and approach This book will be a hands-on guide that will teach you all the what-to-do's, when-to-do's, and how-to-do's for becoming a successful Qlik Sense Consultant. With the help of various business scenarios, the book will cover real-world problems that you can relate to. Various solutions in the book will be backed up by the thought process of why are these solutions used and how you can implement them in your own business environment.

Get hands-on experience in using Burp Suite to execute attacks and perform web assessments Key Features Explore the tools in Burp Suite to meet your web infrastructure security demands Configure Burp to fine-tune the suite of tools specific to the target Use Burp extensions to assist with different technologies commonly found in application stacks Book Description Burp Suite is a Java-based platform for testing the security of your web applications, and has been adopted widely by professional enterprise testers. The Burp Suite Cookbook contains recipes to tackle challenges in determining and exploring vulnerabilities in web applications. You will learn how to uncover security flaws with various test cases for complex environments. After you have configured Burp for your environment, you will use Burp tools such as Spider, Scanner, Intruder, Repeater, and Decoder, among others, to resolve specific problems faced by pentesters. You will also explore working with various modes of Burp and then perform operations on the web. Toward the end, you will cover recipes that target specific test scenarios and resolve them using best practices. By the end of the book, you will be up and running with deploying Burp for securing web applications. What you will learn Configure Burp Suite for your web

applications Perform authentication, authorization, business logic, and data validation testing Explore session management and client-side testing Understand unrestricted file uploads and server-side request forgery Execute XML external entity attacks with Burp Perform remote code execution with Burp Who this book is for If you are a security professional, web pentester, or software developer who wants to adopt Burp Suite for applications security, this book is for you.

IBM® Security Access Manager is a modular, integrated access management appliance that helps secure access to web, mobile, and cloud workloads. It is offered both as a physical appliance and as a virtual appliance image that runs on several popular hypervisors. The integrated appliance form factor enables easier and more flexible deployment and maintenance. This IBM Redpaper™ publication describes the different Security Access Manager Appliance V9.0 deployment patterns and uses hands-on examples to demonstrate how to initially configure systems in those deployments. It also describes various deployment considerations, such as networking, high-availability, performance, disaster recovery, and scalability. All of these deployment patterns are covered within the context of realistic business scenarios. This paper is especially helpful to Security Access Manager architects and deployment specialists.

The Pileri Pioneers profiles twenty-five accomplished BITS Pileri graduates, including startup founders, entrepreneurs and CEOs, who have made an indelible mark on the Indian and global business landscape while deftly navigating the twists and turns of their professional and personal lives. It introduces us to mavericks such as Baba Kalyani, Chairman and Managing Director of Bharat Forge; Hari Menon, Co-founder and CEO of BigBasket; Harish Bhat, Brand Custodian at Tata Sons; Sunil Duggal, Dabur's longest-serving CEO; Tulsi Mirchandaney, MD of Blue Dart Aviation; Nandini Chopra, MD at Alvarez & Marsal; Phanindra Sama, Co-founder of redBus; among others, whose stories of passion and perseverance reveal what it takes to not only achieve success in the competitive arena of business but also make a positive difference in both industry and the world at large. Inspiring, with a touch of nostalgia, this book is an essential read for anyone looking to find their niche and build a meaningful professional life.

Database System Concepts by Silberschatz, Korth and Sudarshan is now in its 6th edition and is one of the cornerstone texts of database education. It presents the fundamental concepts of database management in an intuitive manner geared toward allowing students to begin working with databases as quickly as possible. The text is designed for a first course in databases at the junior/senior undergraduate level or the first year graduate level. It also contains additional material that can be used as supplements or as introductory material for an advanced course. Because the authors present concepts as intuitive descriptions, a familiarity with basic data structures, computer organization, and a high-level programming language are the only prerequisites. Important theoretical results are covered, but formal proofs are omitted. In place of proofs, figures and examples are used to suggest why a result is true.

"It's not easy to find such a generous book on big data and databases. Fortunately, this book is the one." Feng Yu. Computing Reviews. June 28, 2016. This is a book for enterprise architects, database administrators, and developers who need to understand the latest developments in database technologies. It is the book to help you choose the correct database technology at a time when concepts such as Big Data, NoSQL and NewSQL are making what used to be an easy choice into a complex decision with significant implications. The relational database (RDBMS) model completely dominated database technology for over 20 years. Today this "one size fits all" stability has been disrupted by a relatively recent explosion of new database technologies. These paradigm-busting technologies are powering the "Big Data" and "NoSQL" revolutions, as well as forcing fundamental changes in databases across the board. Deciding to use a relational database was once truly a no-brainer, and the various commercial relational databases competed on price, performance, reliability, and ease of use rather than on fundamental architectures. Today we are faced with choices between radically different database technologies. Choosing the right database today is a complex undertaking, with serious economic and technological consequences. Next Generation Databases demystifies today's new database technologies. The book describes what each technology was designed to solve. It shows how each technology can be used to solve real world application and business problems. Most importantly, this book highlights the architectural differences between technologies that are the critical factors to consider when choosing a database platform for new and upcoming projects. Introduces the new technologies that have revolutionized the database landscape Describes how each technology can be used to solve specific application or business challenges Reviews the most popular new wave databases and how they use these new database technologies

If you are a software developer with little or no experience of versioning systems, or are familiar with other centralized versioning systems, then this book is for you. If you have some experience working with command lines or using Linux admin or just using Unix and want to know more about Git, then this book is ideal for you.

BI Publisher for Oracle E-Business Suite is written for Developers and Administrators. It explains how to create BI Publisher reports and how to administer BI Publisher Enterprise. It focuses on your role in an Oracle and Oracle eBS-environment, and explains Concepts, gives Guided exercises, Tips and Exercises. All concepts are explained via real-life examples, screenshots and navigation-paths. Topics: Architecture & installation, Security, XML: XPath, XSL, the Query Builder, Data Templates & Sub-Templates, Conditional coloring, Charts, BIP in the Oracle E-Business Suite, Links to useful websites

In this IBM® Redbooks® publication, we focus on, and provide an overview of, the high availability and Enterprise Replication features of IBM Informix® 11.70. Informix provides solutions for making data highly available in the MACH11 cluster. The components of the MACH11 cluster include High Availability Data Replication (HDR), Shared Disk Secondary (SDS), and Remote Secondary Standby (RSS) servers. Enterprise Replication (ER) provides a means of selectively replicating data between systems in near real time. The Informix Flexible Grid eliminates the administrative complexity of ER. Flexible Grid provides the ability to automatically create database objects, such as tables, indexes, and stored procedures, on all nodes within the grid as a single

operation. These enhanced Enterprise Replication features provide solutions for those customers requiring reliable and quick dissemination of data across a global organization. There is also enhanced capability for customers requiring High Availability Disaster Recovery, in the form of the ability to resend primary server indexes to secondary servers without requiring a rebuild of the index on the primary server. Capabilities such as these enable faster, easier, and more reliable distribution and high availability of data, resulting in improved access and use throughout the enterprise.

Recent advances in the area of lifted inference, which exploits the structure inherent in relational probabilistic models. Statistical relational AI (StaRAI) studies the integration of reasoning under uncertainty with reasoning about individuals and relations. The representations used are often called relational probabilistic models. Lifted inference is about how to exploit the structure inherent in relational probabilistic models, either in the way they are expressed or by extracting structure from observations. This book covers recent significant advances in the area of lifted inference, providing a unifying introduction to this very active field. After providing necessary background on probabilistic graphical models, relational probabilistic models, and learning inside these models, the book turns to lifted inference, first covering exact inference and then approximate inference. In addition, the book considers the theory of liftability and acting in relational domains, which allows the connection of learning and reasoning in relational domains.

Database Systems Using Oracle A Simplified Guide to SQL and PL/SQL Prentice Hall

Learn the art of efficient web scraping and crawling with Python About This Book Extract data from any source to perform real time analytics. Full of techniques and examples to help you crawl websites and extract data within hours. A hands-on guide to web scraping and crawling with real-life problems and solutions Who This Book Is For If you are a software developer, data scientist, NLP or machine-learning enthusiast or just need to migrate your company's wiki from a legacy platform, then this book is for you. It is perfect for someone , who needs instant access to large amounts of semi-structured data effortlessly. What You Will Learn Understand HTML pages and write XPath to extract the data you need Write Scrapy spiders with simple Python and do web crawls Push your data into any database, search engine or analytics system Configure your spider to download files, images and use proxies Create efficient pipelines that shape data in precisely the form you want Use Twisted Asynchronous API to process hundreds of items concurrently Make your crawler super-fast by learning how to tune Scrapy's performance Perform large scale distributed crawls with scrapyd and scrapinghub In Detail This book covers the long awaited Scrapy v 1.0 that empowers you to extract useful data from virtually any source with very little effort. It starts off by explaining the fundamentals of Scrapy framework, followed by a thorough description of how to extract data from any source, clean it up, shape it as per your requirement using Python and 3rd party APIs. Next you will be familiarised with the process of storing the scrapped data in databases as well as search engines and performing real time analytics on them with Spark Streaming. By the end of this book, you will perfect the art of scarping data for your applications with ease Style and approach It is a hands on guide, with first few chapters written as a tutorial, aiming to motivate you and get you started quickly. As the book progresses, more advanced features are explained with real world examples that can be referred while developing your own web applications.

A quick helping guide to every Oracle DBA for understanding interview questions, real time scenario questions.

[Copyright: b655ba8b00f21521028d818d7358515f](https://www.amazon.com/dp/b655ba8b00f21521028d818d7358515f)