

Mongodb The Definitive Guide G C It

As recognized, adventure as with ease as experience virtually lesson, amusement, as with ease as contract can be gotten by just checking out a ebook **mongodb the definitive guide g c it** furthermore it is not directly done, you could say you will even more around this life, with reference to the world.

We find the money for you this proper as competently as easy artifice to acquire those all. We find the money for mongodb the definitive guide g c it and numerous books collections from fictions to scientific research in any way. accompanied by them is this mongodb the definitive guide g c it that can be your partner.

Autonomous Horizons - Greg Zacharias 2019-04-05

Dr. Greg Zacharias, former Chief Scientist of the United States Air Force (2015-18), explores next steps in autonomous systems (AS) development, fielding, and training. Rapid advances in AS development and artificial intelligence (AI) research will change how we think about machines, whether they are individual vehicle platforms or networked enterprises. The payoff will be considerable, affording the US military significant protection for aviators, greater effectiveness in employment, and unlimited opportunities for novel and disruptive concepts of operations. Autonomous Horizons: The Way Forward identifies issues and makes recommendations for the Air Force to take full advantage of this transformational technology.

Principles of Big Graph: In-depth Insight - 2023-02-01

Principles of Big Graph: In-depth Insight, Volume 128 in the Advances in Computer series, highlights new advances in the field with this new volume presenting interesting chapters on a variety of topics, including CESDAM: Centered subgraph data matrix for large graph representation, Bivariate, cluster and suitability analysis of NoSQL Solutions for big graph applications, An empirical investigation on Big Graph using deep learning, Analyzing correlation between quality and accuracy of graph clustering, geneBF: Filtering protein-coded gene

graph data using bloom filter, Processing large graphs with an alternative representation, MapReduce based convolutional graph neural networks: A comprehensive review. Fast exact triangle counting in large graphs using SIMD acceleration, A comprehensive investigation on attack graphs, Qubit representation of a binary tree and its operations in quantum computation, Modified ML-KNN: Role of similarity measures and nearest neighbor configuration in multi label text classification on big social network graph data, Big graph based online learning through social networks, Community detection in large-scale real-world networks, Power rank: An interactive web page ranking algorithm, GA based energy efficient modelling of a wireless sensor network, The major challenges of big graph and their solutions: A review, and An investigation on socio-cyber crime graph. Provides an update on the issues and challenges faced by current researchers Updates on future research agendas Includes advanced topics for intensive research for researchers

Big Data Analytics - Kim H. Pries 2015-02-05

With this book, managers and decision makers are given the tools to make more informed decisions about big data purchasing initiatives. Big Data Analytics: A Practical Guide for Managers not only supplies descriptions of common tools, but also surveys the various products and

vendors that supply the big data market. Comparing and contrasting the dif

PHP and MongoDB Web Development Beginner's Guide - Rubayeet Islam 2011-11-25

Annotation With the rise of Web 2.0, the need for a highly scalable database, capable of storing diverse user-generated content is increasing. MongoDB, an open-source, non-relational database has stepped up to meet this demand and is being used in some of the most popular websites in the world. MongoDB is one of the NoSQL databases which is gaining popularity for developing PHP Web 2.0 applications. PHP and MongoDB Web Development Beginners Guide is a fast-paced, hands-on guide to get started with web application development using PHP and MongoDB. The book follows a Code first, explain later approach, using practical examples in PHP to demonstrate unique features of MongoDB. It does not overwhelm you with information (or starve you of it), but gives you enough to get a solid practical grasp on the concepts. The book starts by introducing the underlying concepts of MongoDB. Each chapter contains practical examples in PHP that teach specific features of the database. The book teaches you to build a blogging application, handle user sessions and authentication, and perform aggregation with MapReduce. You will learn unique MongoDB features and solve interesting problems like real-time analytics, location-aware web apps etc. You will be guided to use MongoDB alongside MySQL to build a diverse data back-end. With its concise coverage of concepts and numerous practical examples, PHP and MongoDB Web Development Beginners Guide is the right choice for the PHP developer to get started with learning MongoDB.

Automated Machine Learning - Frank Hutter 2019-05-17

This open access book presents the first comprehensive overview of general methods in Automated Machine Learning (AutoML), collects descriptions of existing systems based on these methods, and discusses the first series of international challenges of AutoML systems. The recent success of commercial ML applications and the rapid growth of the field has created a high demand for off-the-shelf ML methods that can be used

easily and without expert knowledge. However, many of the recent machine learning successes crucially rely on human experts, who manually select appropriate ML architectures (deep learning architectures or more traditional ML workflows) and their hyperparameters. To overcome this problem, the field of AutoML targets a progressive automation of machine learning, based on principles from optimization and machine learning itself. This book serves as a point of entry into this quickly-developing field for researchers and advanced students alike, as well as providing a reference for practitioners aiming to use AutoML in their work.

Streaming Systems - Tyler Akidau 2018-07-16

Streaming data is a big deal in big data these days. As more and more businesses seek to tame the massive unbounded data sets that pervade our world, streaming systems have finally reached a level of maturity sufficient for mainstream adoption. With this practical guide, data engineers, data scientists, and developers will learn how to work with streaming data in a conceptual and platform-agnostic way. Expanded from Tyler Akidau's popular blog posts "Streaming 101" and "Streaming 102", this book takes you from an introductory level to a nuanced understanding of the what, where, when, and how of processing real-time data streams. You'll also dive deep into watermarks and exactly-once processing with co-authors Slava Chernyak and Reuven Lax. You'll explore: How streaming and batch data processing patterns compare The core principles and concepts behind robust out-of-order data processing How watermarks track progress and completeness in infinite datasets How exactly-once data processing techniques ensure correctness How the concepts of streams and tables form the foundations of both batch and streaming data processing The practical motivations behind a powerful persistent state mechanism, driven by a real-world example How time-varying relations provide a link between stream processing and the world of SQL and relational algebra

Graph Databases - Ian Robinson 2013-06-10

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and

implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems. Learn how different organizations are using graph databases to outperform their competitors. With this book's data modeling, query, and code examples, you'll quickly be able to implement your own solution. Model data with the Cypher query language and property graph model Learn best practices and common pitfalls when modeling with graphs Plan and implement a graph database solution in test-driven fashion Explore real-world examples to learn how and why organizations use a graph database Understand common patterns and components of graph database architecture Use analytical techniques and algorithms to mine graph database information
[Spring Data](#) - Mark Pollack 2012-10-24

You can choose several data access frameworks when building Java enterprise applications that work with relational databases. But what about big data? This hands-on introduction shows you how Spring Data makes it relatively easy to build applications across a wide range of new data access technologies such as NoSQL and Hadoop. Through several sample projects, you'll learn how Spring Data provides a consistent programming model that retains NoSQL-specific features and capabilities, and helps you develop Hadoop applications across a wide range of use-cases such as data analysis, event stream processing, and workflow. You'll also discover the features Spring Data adds to Spring's existing JPA and JDBC support for writing RDBMS-based data access layers. Learn about Spring's template helper classes to simplify the use of database-specific functionality Explore Spring Data's repository abstraction and advanced query functionality Use Spring Data with Redis (key/value store), HBase (column-family), MongoDB (document database), and Neo4j (graph database) Discover the GemFire distributed data grid solution Export Spring Data JPA-managed entities to the Web as RESTful web services Simplify the development of HBase applications, using a lightweight object-mapping framework Build example big-data

pipelines with Spring Batch and Spring Integration

Trino: The Definitive Guide - Matt Fuller 2021-04-14

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

Programming for PaaS - Lucas Carlson 2013-07-31

Platform-as-a-Service (PaaS) is gaining serious traction among web and mobile developers, but as new PaaS providers emerge and existing vendors upgrade their features, it's hard to keep track of what PaaS has to offer. This thorough introduction takes you through the PaaS model from a developer's point of view, and breaks down the types of services that Google App Engine, Windows Azure, Heroku, Cloud Foundry, and others deliver. Whether you're an entrepreneur or part of a large enterprise development team, this book shows you how PaaS can help you focus on innovative applications, rather than spend your time worrying about technical operations. Track the cloud's evolution from IaaS and DevOps to PaaS Learn how PaaS combines the simplicity of shared web hosting with the control of dedicated hosting Explore the benefits of both portable and non-portable PaaS options Apply best practices for moving legacy apps to PaaS—and understand the

challenges involved Write new applications for PaaS from scratch with RESTful meta-services Use PaaS to build mobile apps with backend services that scale Examine the core services that each major provider currently offers Learn the situations in which PaaS might not be advantageous

Data Lake for Enterprises - Tomcy John 2017-05-31

A practical guide to implementing your enterprise data lake using Lambda Architecture as the base About This Book Build a full-fledged data lake for your organization with popular big data technologies using the Lambda architecture as the base Delve into the big data technologies required to meet modern day business strategies A highly practical guide to implementing enterprise data lakes with lots of examples and real-world use-cases Who This Book Is For Java developers and architects who would like to implement a data lake for their enterprise will find this book useful. If you want to get hands-on experience with the Lambda Architecture and big data technologies by implementing a practical solution using these technologies, this book will also help you. What You Will Learn Build an enterprise-level data lake using the relevant big data technologies Understand the core of the Lambda architecture and how to apply it in an enterprise Learn the technical details around Sqoop and its functionalities Integrate Kafka with Hadoop components to acquire enterprise data Use flume with streaming technologies for stream-based processing Understand stream-based processing with reference to Apache Spark Streaming Incorporate Hadoop components and know the advantages they provide for enterprise data lakes Build fast, streaming, and high-performance applications using Elasticsearch Make your data ingestion process consistent across various data formats with configurability Process your data to derive intelligence using machine learning algorithms In Detail The term "Data Lake" has recently emerged as a prominent term in the big data industry. Data scientists can make use of it in deriving meaningful insights that can be used by businesses to redefine or transform the way they operate. Lambda architecture is also emerging as one of the very eminent patterns in the big data landscape, as it not only helps to derive useful information from

historical data but also correlates real-time data to enable business to take critical decisions. This book tries to bring these two important aspects — data lake and lambda architecture—together. This book is divided into three main sections. The first introduces you to the concept of data lakes, the importance of data lakes in enterprises, and getting you up-to-speed with the Lambda architecture. The second section delves into the principal components of building a data lake using the Lambda architecture. It introduces you to popular big data technologies such as Apache Hadoop, Spark, Sqoop, Flume, and Elasticsearch. The third section is a highly practical demonstration of putting it all together, and shows you how an enterprise data lake can be implemented, along with several real-world use-cases. It also shows you how other peripheral components can be added to the lake to make it more efficient. By the end of this book, you will be able to choose the right big data technologies using the lambda architectural patterns to build your enterprise data lake. Style and approach The book takes a pragmatic approach, showing ways to leverage big data technologies and lambda architecture to build an enterprise-level data lake.

CockroachDB - Jesse Seldess 2022

Get the lowdown on CockroachDB, the elastic SQL database built to handle the demands of today's data-driven world. With this practical guide, software developers, architects, and DevOps teams will discover the advantages of building on a distributed SQL database. You'll learn how to create applications that scale elastically and provide seamless delivery for end users while remaining exceptionally resilient and indestructible. Written from scratch for the cloud and architected to scale elastically to handle the demands of cloud native and open source, CockroachDB makes it easier to build and scale modern applications. If you're familiar with distributed systems, you'll quickly discover the benefits of strong data correctness and consistency guarantees as well as optimizations for delivering ultralow latencies to globally distributed end users. With this thorough guide, you'll learn how to: Plan and build applications for distributed infrastructure, including data modeling and schema design Migrate data into CockroachDB Read and write data and

run ACID transactions across distributed infrastructure Optimize queries for performance across geographically distributed replicas Plan a CockroachDB deployment for resiliency across single-region and multiregion clusters Secure, monitor, and optimize your CockroachDB deployment.

NoSQL - Ganesh Chandra Deka 2017-05-19

This book discusses the advanced databases for the cloud-based application known as NoSQL. It will explore the recent advancements in NoSQL database technology. Chapters on structured, unstructured and hybrid databases will be included to explore bigdata analytics, bigdata storage and processing. The book is likely to cover a wide range of topics such as cloud computing, social computing, bigdata and advanced databases processing techniques.

Research and Innovation Forum 2021 - Anna Visvizi 2021

This book features research presented and discussed during the Research and Innovation Forum (Rii Forum) 2021. The Covid-19 pandemic and its social, political, and economic implications had confirmed that a more thorough debate on these issues and topics was needed. For this reason, the Rii Forum 2021 was devoted to the broadly defined question of the short- and long-term impact of the pandemic on our societies. This volume serves as an essential resource to understand the diverse ways in which Covid-19 impacted our societies, including the capacity to innovate, advances in technology, the evolution of the healthcare systems, business model innovation, the prospects of growth, the stability of political systems, and the future of education.

Spring Boot in Action - Craig Walls 2015-12-16

Summary A developer-focused guide to writing applications using Spring Boot. You'll learn how to bypass the tedious configuration steps so that you can concentrate on your application's behavior. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The Spring Framework simplifies enterprise Java development, but it does require lots of tedious configuration work. Spring Boot radically streamlines spinning up a Spring application. You get automatic configuration and a model with

established conventions for build-time and runtime dependencies. You also get a handy command-line interface you can use to write scripts in Groovy. Developers who use Spring Boot often say that they can't imagine going back to hand configuring their applications. About the Book *Spring Boot in Action* is a developer-focused guide to writing applications using Spring Boot. In it, you'll learn how to bypass configuration steps so you can focus on your application's behavior. Spring expert Craig Walls uses interesting and practical examples to teach you both how to use the default settings effectively and how to override and customize Spring Boot for your unique environment. Along the way, you'll pick up insights from Craig's years of Spring development experience. What's Inside Develop Spring apps more efficiently Minimal to no configuration Runtime metrics with the Actuator Covers Spring Boot 1.3 About the Reader Written for readers familiar with the Spring Framework. About the Author Craig Walls is a software developer, author of the popular book *Spring in Action*, Fourth Edition, and a frequent speaker at conferences. Table of Contents Bootstarting Spring Developing your first Spring Boot application Customizing configuration Testing with Spring Boot Getting Groovy with the Spring Boot CLI Applying Grails in Spring Boot Taking a peek inside with the Actuator Deploying Spring Boot applications APPENDIXES Spring Boot developer tools Spring Boot starters Configuration properties Spring Boot dependencies

Java Cookbook - Ian F. Darwin 2020-03-17

Java continues to grow and evolve, and this cookbook continues to evolve in tandem. With this guide, you'll get up to speed right away with hundreds of hands-on recipes across a broad range of Java topics. You'll learn useful techniques for everything from string handling and functional programming to network communication. Each recipe includes self-contained code solutions that you can freely use, along with a discussion of how and why they work. If you're familiar with Java basics, this cookbook will bolster your knowledge of the language and its many recent changes, including how to apply them in your day-to-day development. This updated edition covers changes through Java 12 and

parts of 13 and 14. Recipes include: Methods for compiling, running, and debugging Packaging Java classes and building applications Manipulating, comparing, and rearranging text Regular expressions for string and pattern matching Handling numbers, dates, and times Structuring data with collections, arrays, and other types Object-oriented and functional programming techniques Input/output, directory, and filesystem operations Network programming on both client and server Processing JSON for data interchange Multithreading and concurrency Using Java in big data applications Interfacing Java with other languages

The Way to Go - Ivo Balbaert 2012

This book provides the reader with a comprehensive overview of the new open source programming language Go (in its first stable and maintained release Go 1) from Google. The language is devised with Java / C#-like syntax so as to feel familiar to the bulk of programmers today, but Go code is much cleaner and simpler to read, thus increasing the productivity of developers. You will see how Go: simplifies programming with slices, maps, structs and interfaces incorporates functional programming makes error-handling easy and secure simplifies concurrent and parallel programming with goroutines and channels And you will learn how to: make use of Go's excellent standard library program Go the idiomatic way using patterns and best practices in over 225 working examples and 135 exercises This book focuses on the aspects that the reader needs to take part in the coming software revolution using Go.

Hands-On System Programming with Go - Alex Guerrieri 2019-07-05

Explore the fundamentals of systems programming starting from kernel API and filesystem to network programming and process communications Key Features Learn how to write Unix and Linux system code in Golang v1.12 Perform inter-process communication using pipes, message queues, shared memory, and semaphores Explore modern Go features such as goroutines and channels that facilitate systems programming Book Description System software and applications were largely created using low-level languages such as C or C++. Go is a modern language that combines simplicity, concurrency, and

performance, making it a good alternative for building system applications for Linux and macOS. This Go book introduces Unix and systems programming to help you understand the components the OS has to offer, ranging from the kernel API to the filesystem, and familiarize yourself with Go and its specifications. You'll also learn how to optimize input and output operations with files and streams of data, which are useful tools in building pseudo terminal applications. You'll gain insights into how processes communicate with each other, and learn about processes and daemon control using signals, pipes, and exit codes. This book will also enable you to understand how to use network communication using various protocols, including TCP and HTTP. As you advance, you'll focus on Go's best feature-concurrency helping you handle communication with channels and goroutines, other concurrency tools to synchronize shared resources, and the context package to write elegant applications. By the end of this book, you will have learned how to build concurrent system applications using Go What you will learn Explore concepts of system programming using Go and concurrency Gain insights into Golang's internals, memory models and allocation Familiarize yourself with the filesystem and IO streams in general Handle and control processes and daemons' lifetime via signals and pipes Communicate with other applications effectively using a network Use various encoding formats to serialize complex data structures Become well-versed in concurrency with channels, goroutines, and sync Use concurrency patterns to build robust and performant system applications Who this book is for If you are a developer who wants to learn system programming with Go, this book is for you. Although no knowledge of Unix and Linux system programming is necessary, intermediate knowledge of Go will help you understand the concepts covered in the book

The JHipster Mini-Book - Matt Raible 2016-11-25

The things you need to do to set up a new software project can be daunting. First, you have to select the back-end framework to create your API, choose your database, set up security, and choose your build tool. Then you have to choose the tools to create your front end: select a

UI framework, configure a build tool, set up Sass processing, configure your browser to auto-refresh when you make changes, and configure the client and server so they work in unison. If you're building a new application using Spring Boot and Angular, you can save days by using JHipster. JHipster generates a complete and modern web app, unifying: - A high-performance and robust Java stack on the server side with Spring Boot - A sleek, modern, mobile-first front-end with Angular and Bootstrap - A robust microservice architecture with the JHipster Registry, Netflix OSS, the ELK stack, and Docker - A powerful workflow to build your application with Yeoman, Webpack, and Maven/Gradle

Kubernetes Operators - Jason Dobies 2020-02-21

Operators are a way of packaging, deploying, and managing Kubernetes applications. A Kubernetes application doesn't just run on Kubernetes; it's composed and managed in Kubernetes terms. Operators add application-specific operational knowledge to a Kubernetes cluster, making it easier to automate complex, stateful applications and to augment the platform. Operators can coordinate application upgrades seamlessly, react to failures automatically, and streamline repetitive maintenance like backups. Think of Operators as site reliability engineers in software. They work by extending the Kubernetes control plane and API, helping systems integrators, cluster administrators, and application developers reliably deploy and manage key services and components. Using real-world examples, authors Jason Dobies and Joshua Wood demonstrate how to use Operators today and how to create Operators for your applications with the Operator Framework and SDK. Learn how to establish a Kubernetes cluster and deploy an Operator Examine a range of Operators from usage to implementation Explore the three pillars of the Operator Framework: the Operator SDK, the Operator Lifecycle Manager, and Operator Metering Build Operators from the ground up using the Operator SDK Build, package, and run an Operator in development, testing, and production phases Learn how to distribute your Operator for installation on Kubernetes clusters

Cassandra: The Definitive Guide - Jeff Carpenter 2016-06-29

Imagine what you could do if scalability wasn't a problem. With this

hands-on guide, you'll learn how the Cassandra database management system handles hundreds of terabytes of data while remaining highly available across multiple data centers. This expanded second edition—updated for Cassandra 3.0—provides the technical details and practical examples you need to put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's non-relational design, with special attention to data modeling. If you're a developer, DBA, or application architect looking to solve a database scaling issue or future-proof your application, this guide helps you harness Cassandra's speed and flexibility.

Understand Cassandra's distributed and decentralized structure Use the Cassandra Query Language (CQL) and cqlsh—the CQL shell Create a working data model and compare it with an equivalent relational model Develop sample applications using client drivers for languages including Java, Python, and Node.js Explore cluster topology and learn how nodes exchange data Maintain a high level of performance in your cluster Deploy Cassandra on site, in the Cloud, or with Docker Integrate Cassandra with Spark, Hadoop, Elasticsearch, Solr, and Lucene [Smart Devices, Applications, and Protocols for the IoT](#) - Rodrigues, Joel J. P. C. 2019-03-29

Advances in computing, communications, and control have bridged the physical components of reality and cyberspace leading to the smart internet of things (IoT). The notion of IoT has extraordinary significance for the future of several industrial domains. Hence, it is expected that the complexity in the design of IoT applications will continue to increase due to the integration of several cyber components with physical and industrial systems. As a result, several smart protocols and algorithms are needed to communicate and exchange data between IoT devices. Smart Devices, Applications, and Protocols for the IoT is a collection of innovative research that explores new methods and techniques for achieving reliable and efficient communication in recent applications including machine learning, network optimization, adaptive methods, and smart algorithms and protocols. While highlighting topics including artificial intelligence, sensor networks, and mobile network

architectures, this book is ideally designed for IT specialists and consultants, software engineers, technology developers, academicians, researchers, and students seeking current research on up-to-date technologies in smart communications, protocols, and algorithms in IoT.

Java 11 Cookbook - Nick Samoylov 2018-09-29

Solutions for modular, functional, reactive, GUI, network, and multithreaded programming Key Features Explore the latest features of Java 11 to implement efficient and reliable code Develop memory-efficient applications, understanding new garbage collection in Java 11 Create restful webservices and microservices with Spring boot 2 and Docker Book Description For more than three decades, Java has been on the forefront of developing robust software that has helped versatile businesses meet their requirements. Being one of the most widely used programming languages in history, it's imperative for Java developers to discover effective ways of using it in order to take full advantage of the power of the latest Java features. Java 11 Cookbook offers a range of software development solutions with simple and straightforward Java 11 code examples to help you build a modern software system. Starting with the installation of Java, each recipe addresses various problem by explaining the solution and offering insights into how it works. You'll explore the new features added to Java 11 that will make your application modular, secure, and fast. The book contains recipes on functional programming, GUI programming, concurrent programming, and database programming in Java. You'll also be taken through the new features introduced in JDK 11.3 and 11.9. By the end of this book, you'll be equipped with the skills required to write robust, scalable, and optimal Java code effectively. What you will learn Set up JDK and understand what's new in the JDK 11 installation Implement object-oriented designs using classes and interfaces Manage operating system processes Create a modular application with clear dependencies Build graphical user interfaces using JavaFX Use the new HTTP Client API Explore the new diagnostic features in Java 11 Discover how to use the new JShell REPL tool Who this book is for The book is for intermediate-to-advanced Java programmers who want to make their

applications fast, secure, and scalable.

Kafka: The Definitive Guide - Neha Narkhede 2017-08-31

Every enterprise application creates data, whether it's log messages, metrics, user activity, outgoing messages, or something else. And how to move all of this data becomes nearly as important as the data itself. If you're an application architect, developer, or production engineer new to Apache Kafka, this practical guide shows you how to use this open source streaming platform to handle real-time data feeds. Engineers from Confluent and LinkedIn who are responsible for developing Kafka explain how to deploy production Kafka clusters, write reliable event-driven microservices, and build scalable stream-processing applications with this platform. Through detailed examples, you'll learn Kafka's design principles, reliability guarantees, key APIs, and architecture details, including the replication protocol, the controller, and the storage layer. Understand publish-subscribe messaging and how it fits in the big data ecosystem. Explore Kafka producers and consumers for writing and reading messages Understand Kafka patterns and use-case requirements to ensure reliable data delivery Get best practices for building data pipelines and applications with Kafka Manage Kafka in production, and learn to perform monitoring, tuning, and maintenance tasks Learn the most critical metrics among Kafka's operational measurements Explore how Kafka's stream delivery capabilities make it a perfect source for stream processing systems

Big Data Beyond the Hype - Zikopoulos 2014-11-10

Big Data in a nutshell: It is the ability to retain, process, and understand data like never before. It can mean more data than what you are using today; but it can also mean different kinds of data, a venture into the unstructured world where most of today's data resides. In this book you will learn how cognitive computing systems, like IBM Watson, fit into the Big Data world. Learn about the concept of data-in-motion and InfoSphere Streams, the world's fastest and most flexible platform for streaming data. Capturing, storing, refining, transforming, governing, securing, and analyzing data are important topics also covered in this book.

Microsoft Visual Studio 2015 Unleashed - Lars Powers 2015-08-07

Microsoft Visual Studio 2015 empowers you to write next-generation applications for any modern environment: mobile, web, cloud, universal Windows 10/8.x, database, and beyond. This end-to-end deep dive will help working developers squeeze maximum productivity out of Microsoft's powerful new toolset. The authors combine authoritative and detailed information about Microsoft's latest IDE, with extensive insights and best practices drawn from decades of development experience. Developers will quickly get comfortable with Visual Studio 2015's updated interface, master its new capabilities, leverage its extensive new support for open standards, and discover multiple opportunities to leverage its .NET 4.6 platform and language improvements. By focusing entirely on Visual Studio 2015 Professional, the authors go deeper into Microsoft's core product than ever before. You'll find expert coverage of everything from debugging through deploying to Azure, IDE extension and automation through cross-platform mobile development. Throughout, this book's focus is relentlessly practical: how to apply Microsoft's tools to build better software, faster. Detailed information on how to... Master Visual Studio 2015's updated interface and key tools: Solutions, Projects, Browsers, Explorers, Editors, and Designers to improve productivity Develop robust cross-platform mobile apps for Windows, iOS, and Android using Apache Cordova templates for Visual Studio Use the new ASP.NET 5 to build modern web solutions that run on Windows, Mac, or Linux Develop Single Page Applications (SPAs) based on HTML5 and rich client-side JavaScript frameworks such as Knockout, AngularJS, Bootstrap, and more Accelerate cloud development with the Azure SDK, QuickStart templates, and Azure management portal Create mobile service solutions using ASP.NET Web API and WCF Streamline data development across multiple platforms with Entity Framework 7 Develop modern Microsoft Office business applications Perform robust, automated unit testing as you code, increasing your confidence in changes and refactoring Extend the VS 2015 IDE and Code Editor by creating custom, productivity-enhancing solutions Download all examples and source code presented in this book from

informit.com/title/9780672337369 as they become available.

SQL Server 2016 Developer's Guide - Dejan Sarka 2017-03-22

Get the most out of the rich development capabilities of SQL Server 2016 to build efficient database applications for your organization About This Book Utilize the new enhancements in Transact-SQL and security features in SQL Server 2016 to build efficient database applications Work with temporal tables to get information about data stored in the table at any point in time A detailed guide to SQL Server 2016, introducing you to multiple new features and enhancements to improve your overall development experience Who This Book Is For This book is for database developers and solution architects who plan to use the new SQL Server 2016 features for developing efficient database applications. It is also ideal for experienced SQL Server developers who want to switch to SQL Server 2016 for its rich development capabilities. Some understanding of the basic database concepts and Transact-SQL language is assumed. What You Will Learn Explore the new development features introduced in SQL Server 2016 Identify opportunities for In-Memory OLTP technology, significantly enhanced in SQL Server 2016 Use columnstore indexes to get significant storage and performance improvements Extend database design solutions using temporal tables Exchange JSON data between applications and SQL Server in a more efficient way Migrate historical data transparently and securely to Microsoft Azure by using Stretch Database Use the new security features to encrypt or to have more granular control over access to rows in a table Simplify performance troubleshooting with Query Store Discover the potential of R's integration with SQL Server In Detail Microsoft SQL Server 2016 is considered the biggest leap in the data platform history of the Microsoft, in the ongoing era of Big Data and data science. Compared to its predecessors, SQL Server 2016 offers developers a unique opportunity to leverage the advanced features and build applications that are robust, scalable, and easy to administer. This book introduces you to new features of SQL Server 2016 which will open a completely new set of possibilities for you as a developer. It prepares you for the more advanced topics by starting with a quick introduction to

SQL Server 2016's new features and a recapitulation of the possibilities you may have already explored with previous versions of SQL Server. The next part introduces you to small delights in the Transact-SQL language and then switches to a completely new technology inside SQL Server - JSON support. We also take a look at the Stretch database, security enhancements, and temporal tables. The last chapters concentrate on implementing advanced topics, including Query Store, columnstore indexes, and In-Memory OLTP. You will finally be introduced to R and how to use the R language with Transact-SQL for data exploration and analysis. By the end of this book, you will have the required information to design efficient, high-performance database applications without any hassle. Style and approach This book is a detailed guide to mastering the development features offered by SQL Server 2016, with a unique learn-as-you-do approach. All the concepts are explained in a very easy-to-understand manner and are supplemented with examples to ensure that you—the developer—are able to take that next step in building more powerful, robust applications for your organization with ease.

Professional NoSQL - Shashank Tiwari 2011-08-31

A hands-on guide to leveraging NoSQL databases NoSQL databases are an efficient and powerful tool for storing and manipulating vast quantities of data. Most NoSQL databases scale well as data grows. In addition, they are often malleable and flexible enough to accommodate semi-structured and sparse data sets. This comprehensive hands-on guide presents fundamental concepts and practical solutions for getting you ready to use NoSQL databases. Expert author Shashank Tiwari begins with a helpful introduction on the subject of NoSQL, explains its characteristics and typical uses, and looks at where it fits in the application stack. Unique insights help you choose which NoSQL solutions are best for solving your specific data storage needs.

Professional NoSQL: Demystifies the concepts that relate to NoSQL databases, including column-family oriented stores, key/value databases, and document databases. Delves into installing and configuring a number of NoSQL products and the Hadoop family of products. Explains

ways of storing, accessing, and querying data in NoSQL databases through examples that use MongoDB, HBase, Cassandra, Redis, CouchDB, Google App Engine Datastore and more. Looks at architecture and internals. Provides guidelines for optimal usage, performance tuning, and scalable configurations. Presents a number of tools and utilities relating to NoSQL, distributed platforms, and scalable processing, including Hive, Pig, RRDtool, Nagios, and more.

Learning Spark - Jules S. Damji 2020-07-16

Data is bigger, arrives faster, and comes in a variety of formats—and it all needs to be processed at scale for analytics or machine learning. But how can you process such varied workloads efficiently? Enter Apache Spark. Updated to include Spark 3.0, this second edition shows data engineers and data scientists why structure and unification in Spark matters. Specifically, this book explains how to perform simple and complex data analytics and employ machine learning algorithms. Through step-by-step walk-throughs, code snippets, and notebooks, you'll be able to: Learn Python, SQL, Scala, or Java high-level Structured APIs Understand Spark operations and SQL Engine Inspect, tune, and debug Spark operations with Spark configurations and Spark UI Connect to data sources: JSON, Parquet, CSV, Avro, ORC, Hive, S3, or Kafka Perform analytics on batch and streaming data using Structured Streaming Build reliable data pipelines with open source Delta Lake and Spark Develop machine learning pipelines with MLlib and productionize models using MLflow

Spark: The Definitive Guide - Bill Chambers 2018-02-08

Learn how to use, deploy, and maintain Apache Spark with this comprehensive guide, written by the creators of the open-source cluster-computing framework. With an emphasis on improvements and new features in Spark 2.0, authors Bill Chambers and Matei Zaharia break down Spark topics into distinct sections, each with unique goals. You'll explore the basic operations and common functions of Spark's structured APIs, as well as Structured Streaming, a new high-level API for building end-to-end streaming applications. Developers and system administrators will learn the fundamentals of monitoring, tuning, and debugging Spark,

and explore machine learning techniques and scenarios for employing MLlib, Spark's scalable machine-learning library. Get a gentle overview of big data and Spark Learn about DataFrames, SQL, and Datasets—Spark's core APIs—through worked examples Dive into Spark's low-level APIs, RDDs, and execution of SQL and DataFrames Understand how Spark runs on a cluster Debug, monitor, and tune Spark clusters and applications Learn the power of Structured Streaming, Spark's stream-processing engine Learn how you can apply MLlib to a variety of problems, including classification or recommendation

Engineering Production-Grade Shiny Apps - Colin Fay 2021-09-27

From the Reviews "[This book] contains an excellent blend of both Shiny-specific topics ... and practical advice from software development that fits in nicely with Shiny apps. You will find many nuggets of wisdom sprinkled throughout these chapters...." Eric Nantz, Host of the R-Podcast and the Shiny Developer Series (from the Foreword) "[This] book is a gradual and pleasant invitation to the production-ready shiny apps world. It ...exposes a comprehensive and robust workflow powered by the {golem} package. [It] fills the not yet covered gap between shiny app development and deployment in such a thrilling way that it may be read in one sitting.... In the industry world, where processes robustness is a key toward productivity, this book will indubitably have a tremendous impact." David Granjon, Sr. Expert Data Science, Novartis Presented in full color, *Engineering Production-Grade Shiny Apps* helps people build production-grade shiny applications, by providing advice, tools, and a methodology to work on web applications with R. This book starts with an overview of the challenges which arise from any big web application project: organizing work, thinking about the user interface, the challenges of teamwork and the production environment. Then, it moves to a step-by-step methodology that goes from the idea to the end application. Each part of this process will cover in detail a series of tools and methods to use while building production-ready shiny applications. Finally, the book will end with a series of approaches and advice about optimizations for production. Features Focused on practical matters: This book does not cover Shiny concepts, but practical tools and

methodologies to use for production. Based on experience: This book is a formalization of several years of experience building Shiny applications. Original content: This book presents new methodologies and tooling, not just a review of what already exists. *Engineering Production-Grade Shiny Apps* covers medium to advanced content about Shiny, so it will help people that are already familiar with building apps with Shiny, and who want to go one step further.

Node.js, MongoDB and Angular Web Development - Brad Dayley
2017-10-11

Node.js, MongoDB and Angular Web Development The definitive guide to using the MEAN stack to build web applications Node.js is a leading server-side programming environment, MongoDB is the most popular NoSQL database, and Angular is the leading framework for MVC-based front-end development. Together, they provide an easy-to-implement, fully integrated web development stack that allows web programmers to create high-performance sites and applications built completely in JavaScript, from server to client. Updated for Angular 2, Angular 4, and subsequent versions, this new edition of *Node.js, MongoDB and Angular Web Development* shows you how to integrate these three technologies into complete working solutions. It begins with concise, crystal-clear tutorials on each technology and then quickly moves on to building common web applications. You'll learn how to use Node.js and MongoDB to build more scalable, high-performance sites, how to leverage Angular's innovative MVC approach to structure more effective pages and applications, and how to use all three together to deliver outstanding next-generation Web solutions. Implement a highly scalable and dynamic web server using Node.js and Express Implement a MongoDB data store for your web applications Access and interact with MongoDB from Node.js JavaScript code Learn the basics of TypeScript Define custom Angular directives that extend the HTML language Build server-side web services in JavaScript Implement client-side services that can interact with the Node.js web server Build dynamic browser views that provide rich user interaction Add authenticated user accounts and nested comment components to your web applications and pages Contents at a

Glance Part I: Getting Started 1 Introducing the Node.js-to-Angular Stack 2 JavaScript Primer Part II: Learning Node.js 3 Getting Started with Node.js 4 Using Events, Listeners, Timers, and Callbacks in Node.js 5 Handling Data I/O in Node.js 6 Accessing the File System from Node.js 7 Implementing HTTP Services in Node.js 8 Implementing Socket Services in Node.js 9 Scaling Applications Using Multiple Processors in Node.js 10 Using Additional Node.js Modules Part III: Learning MongoDB 11 Understanding NoSQL and MongoDB 12 Getting Started with MongoDB 13 Getting Started with MongoDB and Node.js 14 Manipulating MongoDB Documents from Node.js 15 Accessing MongoDB from Node.js 16 Using Mongoose for Structured Schema and Validation 17 Advanced MongoDB Concepts Part IV: Using Express to Make Life Easier 18 Implementing Express in Node.js 19 Implementing Express Middleware Part V: Learning Angular 20 Jumping into TypeScript 21 Getting Started with Angular 22 Angular Components 23 Expressions 24 Data Binding 25 Built-in Directives Part VI: Advanced Angular 26 Custom Directives 27 Events and Change Detection 28 Implementing Angular Services in Web Applications 29 Creating Your Own Custom Angular Services 30 Having Fun with Angular

Graph Databases - Ian Robinson 2015-06-10

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems. This second edition includes new code samples and diagrams, using the latest Neo4j syntax, as well as information on new functionality. Learn how different organizations are using graph databases to outperform their competitors. With this book's data modeling, query, and code examples, you'll quickly be able to implement your own solution. Model data with the Cypher query language and property graph model Learn best practices and common pitfalls when modeling with graphs Plan and implement a graph database

solution in test-driven fashion Explore real-world examples to learn how and why organizations use a graph database Understand common patterns and components of graph database architecture Use analytical techniques and algorithms to mine graph database information *Seven Databases in Seven Weeks* - Luc Perkins 2018-04-05 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.

Beginning Apache Pig - Balaswamy Vaddeman 2016-12-10

Learn to use Apache Pig to develop lightweight big data applications easily and quickly. This book shows you many optimization techniques and covers every context where Pig is used in big data analytics. Beginning Apache Pig shows you how Pig is easy to learn and requires relatively little time to develop big data applications. The book is divided into four parts: the complete features of Apache Pig; integration with other tools; how to solve complex business problems; and optimization of tools. You'll discover topics such as MapReduce and why it cannot meet every business need; the features of Pig Latin such as data types for each load, store, joins, groups, and ordering; how Pig workflows can be created; submitting Pig jobs using Hue; and working with Oozie. You'll also see how to extend the framework by writing UDFs and custom load, store, and filter functions. Finally you'll cover different optimization techniques such as gathering statistics about a Pig script, joining strategies, parallelism, and the role of data formats in good performance. What You Will Learn • Use all the features of Apache Pig • Integrate Apache Pig with other tools • Extend Apache Pig • Optimize Pig Latin code • Solve different use cases for Pig Latin Who This Book Is For All levels of IT professionals: architects, big data enthusiasts, engineers, developers, and big data administrators

Cloud Native Go - Kevin Hoffman 2016-12-05

The Complete Guide to Building Cloud-Based Services Cloud Native Go shows developers how to build massive cloud applications that meet the insatiable demands of today's customers, and will dynamically scale to handle virtually any volume of data, traffic, or users. Kevin Hoffman and Dan Nemeth describe the modern cloud-native application in detail, illuminating factors, disciplines, and habits associated with rapid, reliable cloud-native development. They also introduce Go, a "simply elegant" high-performance language that is especially well-suited for cloud development. You'll walk through creating microservices in Go, adding front-end web components using ReactJS and Flux, and mastering advanced Go-based cloud-native techniques. Hoffman and Nemeth show

how to build a continuous delivery pipeline with tools like Wercker, Docker, and Dockerhub; automatically push apps to leading platforms; and systematically monitor app performance in production. Learn "The Way of the Cloud": why developing good cloud software is fundamentally about mindset and discipline Discover why Go is ideal for cloud-native microservices development Plan cloud apps that support continuous delivery and deployment Design service ecosystems, and then build them in a test-first manner Push work-in-progress to a cloud Use Event Sourcing and CQRS patterns to react and respond to enormous volume and throughput Secure cloud-based web applications: do's, don'ts, and options Create reactive applications in the cloud with third-party messaging providers Build massive-scale, cloud-friendly GUIs with React and Flux Monitor dynamic scaling, failover, and fault tolerance in the cloud

Oracle on LinuxONE - Susan Adamovich 2019-08-08

Oracle Database 12c running on Linux is available for deployment on IBM® LinuxONE. The enterprise-grade Linux on LinuxONE solution is designed to add value to Oracle Database solutions, including the new functions that are introduced in Oracle Database 12c. In this IBM Redbooks® publication, we explore the IBM and Oracle Alliance and describe how Oracle Database benefits from LinuxONE. We then explain how to set up Linux guests to install Oracle Database 12c. We also describe how to use the Oracle Enterprise Manager Cloud Control Agent to manage Oracle Database 12c Release 1. Additionally, we discuss encryption for Oracle using Oracle Transparent Data Encryption (TDE) on Oracle 12c Release 2. We also describe a successful consolidation project from sizing to migration, performance management topics, and high availability. Finally, we end with a chapter about surrounding Oracle with Open Source software. The audience for this publication includes database consultants, installers, administrators, and system programmers. This publication is not meant to replace Oracle documentation, but to supplement it with our experiences while installing and using Oracle products.

N1ql - Sitaram Vemulapalli 2018-04-18

The Definitive Guide to Spring Batch - Michael T. Minella 2019-07-08
Work with all aspects of batch processing in a modern Java environment using a selection of Spring frameworks. This book provides up-to-date examples using the latest configuration techniques based on Java configuration and Spring Boot. The Definitive Guide to Spring Batch takes you from the “Hello, World!” of batch processing to complex scenarios demonstrating cloud native techniques for developing batch applications to be run on modern platforms. Finally this book demonstrates how you can use areas of the Spring portfolio beyond just Spring Batch 4 to collaboratively develop mission-critical batch processes. You’ll see how a new class of use cases and platforms has evolved to have an impact on batch-processing. Data science and big data have become prominent in modern IT and the use of batch processing to orchestrate workloads has become commonplace. The Definitive Guide to Spring Batch covers how running finite tasks on cloud infrastructure in a standardized way has changed where batch applications are run. Additionally, you’ll discover how Spring Batch 4 takes advantage of Java 9, Spring Framework 5, and the new Spring Boot 2 micro-framework. After reading this book, you’ll be able to use Spring Boot to simplify the development of your own Spring projects, as well as take advantage of Spring Cloud Task and Spring Cloud Data Flow for added cloud native functionality. Includes a foreword by Dave Syer,

Spring Batch project founder. What You'll Learn Discover what is new in Spring Batch 4 Carry out finite batch processing in the cloud using the Spring Batch project Understand the newest configuration techniques based on Java configuration and Spring Boot using practical examples Master batch processing in complex scenarios including in the cloud Develop batch applications to be run on modern platforms Use areas of the Spring portfolio beyond Spring Batch to develop mission-critical batch processes Who This Book Is For Experienced Java and Spring coders new to the Spring Batch platform. This definitive book will be useful in allowing even experienced Spring Batch users and developers to maximize the Spring Batch tool.

Handbook of Research on Cloud Infrastructures for Big Data Analytics - Raj, Pethuru 2014-03-31

Clouds are being positioned as the next-generation consolidated, centralized, yet federated IT infrastructure for hosting all kinds of IT platforms and for deploying, maintaining, and managing a wider variety of personal, as well as professional applications and services. Handbook of Research on Cloud Infrastructures for Big Data Analytics focuses exclusively on the topic of cloud-sponsored big data analytics for creating flexible and futuristic organizations. This book helps researchers and practitioners, as well as business entrepreneurs, to make informed decisions and consider appropriate action to simplify and streamline the arduous journey towards smarter enterprises.