



databases, using the right compression and compaction strategies, and loading and unloading data. Expert Apache Cassandra Administration provides numerous step-by-step examples starting with the basics of a Cassandra database, and going all the way through backup and recovery, performance optimization, and monitoring and securing the data. The book serves as an authoritative and comprehensive guide to the building and management of simple to complex Cassandra databases. The book:

- Takes you through building a Cassandra database from installation of the software and creation of a single database, through to complex clusters and data centers
- Provides numerous examples of actual commands in a real-life Cassandra environment that show how to confidently configure, manage, troubleshoot, and tune Cassandra databases
- Shows how to use the Cassandra configuration properties to build a highly stable, available, and secure Cassandra database that always operates at peak efficiency

What You'll Learn

- Install the Cassandra software and create your first database
- Understand the Cassandra data model, and the internal architecture of a Cassandra database
- Create your own Cassandra cluster, step-by-step
- Run a Cassandra cluster on Docker
- Work with Apache Spark by connecting to a Cassandra database
- Deploy Cassandra clusters in your data center, or on Amazon EC2 instances
- Back up and restore mission-critical Cassandra databases
- Monitor, troubleshoot, and tune production Cassandra databases, and cut your spending on resources such as memory, servers, and storage

Who This Book Is For

Database administrators, developers, and architects who are looking for an authoritative and comprehensive single volume for all their Cassandra administration needs. Also for administrators who are tasked with setting up and maintaining highly reliable and high-performing Cassandra databases. An excellent choice for big data administrators, database administrators, architects, and developers who use Cassandra as their key data store, to support high volume online transactions, or as a decentralized, elastic data store.

Over 80 recipes to simplify machine learning model implementations with Spark

About This Book

- Solve the day-to-day problems of data science with Spark
- This unique cookbook consists of exciting and intuitive numerical recipes
- Optimize your work by acquiring, cleaning, analyzing, predicting, and visualizing your data

Who This Book Is For

This book is for Scala developers with a fairly good exposure to and understanding of machine learning techniques, but lack practical implementations with Spark. A solid knowledge of machine learning algorithms is assumed, as well as hands-on experience of implementing ML algorithms with Scala. However, you do not need to be acquainted with the Spark ML libraries and ecosystem.

What You Will Learn

- Get to know how Scala and Spark go hand-in-hand for developers when developing ML systems with Spark
- Build a recommendation engine that scales with Spark
- Find out how to build unsupervised clustering systems to classify data in Spark
- Build machine learning systems with the Decision Tree and Ensemble models in Spark
- Deal with the curse of high-dimensionality in big data using Spark
- Implement Text analytics for Search Engines in Spark
- Streaming Machine Learning System implementation using Spark

In Detail

Machine learning aims to extract knowledge from data, relying on fundamental concepts in computer science, statistics, probability, and optimization. Learning about algorithms enables a wide range of applications, from everyday tasks such as product recommendations and spam filtering to bleeding edge applications such as self-driving cars and personalized medicine. You will gain hands-on experience of applying these principles using Apache Spark, a cluster computing system well suited for large-scale machine learning tasks. This book begins with a quick overview of setting up the necessary IDEs to facilitate the execution of code

examples that will be covered. It also highlights some key issues developers face while thinking about Scala for machine learning and during the switch over to Spark. We progress by uncovering the various Spark APIs and the implementation of ML algorithms with developing classification systems, recommendation engines, clustering and learning systems. Towards the final chapters, we'll focus on building high-end applications and explain various unsupervised methodologies and challenges to tackle when implementing with big data ML systems. This practical guide explains you to program and understand the power of Apache Cassandra 3.x. You will explore the integration and interaction of Cassandra components, and explore features such as the token allocation algorithm, CQL3, vnodes, lightweight transactions, and data modelling in detail.

The ultimate guide to managing, building, and deploying large-scale clusters with Apache Mesos About This Book Master the architecture of Mesos and intelligently distribute your task across clusters of machines Explore a wide range of tools and platforms that Mesos works with This real-world comprehensive and robust tutorial will help you become an expert Who This Book Is For The book aims to serve DevOps engineers and system administrators who are familiar with the basics of managing a Linux system and its tools What You Will Learn Understand the Mesos architecture Manually spin up a Mesos cluster on a distributed infrastructure Deploy a multi-node Mesos cluster using your favorite DevOps See the nuts and bolts of scheduling, service discovery, failure handling, security, monitoring, and debugging in an enterprise-grade, production cluster deployment Use Mesos to deploy big data frameworks, containerized applications, or even custom build your own applications effortlessly In Detail Apache Mesos is open source cluster management software that provides efficient resource isolations and resource sharing distributed applications or frameworks. This book will take you on a journey to enhance your knowledge from amateur to master level, showing you how to improve the efficiency, management, and development of Mesos clusters. The architecture is quite complex and this book will explore the difficulties and complexities of working with Mesos. We begin by introducing Mesos, explaining its architecture and functionality. Next, we provide a comprehensive overview of Mesos features and advanced topics such as high availability, fault tolerance, scaling, and efficiency. Furthermore, you will learn to set up multi-node Mesos clusters on private and public clouds. We will also introduce several Mesos-based scheduling and management frameworks or applications to enable the easy deployment, discovery, load balancing, and failure handling of long-running services. Next, you will find out how a Mesos cluster can be easily set up and monitored using the standard deployment and configuration management tools. This advanced guide will show you how to deploy important big data processing frameworks such as Hadoop, Spark, and Storm on Mesos and big data storage frameworks such as Cassandra, Elasticsearch, and Kafka. Style and approach This advanced guide provides a detailed step-by-step account of deploying a Mesos cluster. It will demystify the concepts behind Mesos.

This book will be helpful for those who is preparing for interview or getting interviewed. It is specially designed to brush-up the java concepts quickly.

This book constitutes the refereed proceedings of the 22 International Conference on Database and Expert Systems Applications, DEXA 2011, held in Toulouse, France, August 29 - September 2, 2011. The 52 revised full papers and 40 short papers presented were carefully reviewed and selected from 207 submissions. The papers are organized in topical sections on query processing; database semantics; skyline queries; security and privacy; spatial and temporal data; semantic web search; storage and search; web search; data integration, transactions and optimization; and web applications.

????????????????????????????, SQL,?????,????????????????????????????

Build efficient data flow and machine learning programs with this flexible, multi-functional open-source cluster-computing framework Key Features Master the art of real-time big data processing and machine learning Explore a wide range of use-cases to analyze large data Discover ways to optimize your work by using many features of Spark 2.x and Scala Book Description Apache Spark is an in-memory, cluster-based data processing system that provides a wide range of functionalities such as big data processing, analytics, machine learning, and more. With this Learning Path, you can take your knowledge of Apache Spark to the next level by learning how to expand Spark's functionality and building your own data flow and machine learning programs on this platform. You will work with the different modules in Apache Spark, such as interactive querying with Spark SQL, using DataFrames and datasets, implementing streaming analytics with Spark Streaming, and applying machine learning and deep learning techniques on Spark using MLlib and various external tools. By the end of this elaborately designed Learning Path, you will have all the knowledge you need to master Apache Spark, and build your own big data processing and analytics pipeline quickly and without any hassle. This Learning Path includes content from the following Packt products: Mastering Apache Spark 2.x by Romeo Kienzler Scala and Spark for Big Data Analytics by Md. Rezaul Karim, Sridhar Alla Apache Spark 2.x Machine Learning Cookbook by Siamak Amirghodsi, Meenakshi Rajendran, Broderick Hall, Shuen Mei Cookbook What you will learn Get to grips with all the features of Apache Spark 2.x Perform highly optimized real-time big data processing Use ML and DL techniques with Spark MLlib and third-party tools Analyze structured and unstructured data using SparkSQL and GraphX Understand tuning, debugging, and monitoring of big data applications Build scalable and fault-tolerant streaming applications Develop scalable recommendation engines Who this book is for If you are an intermediate-level Spark developer looking to master the advanced capabilities and use-cases of Apache Spark 2.x, this Learning Path is ideal for you. Big data professionals who want to learn how to integrate and use the features of Apache Spark and build a strong big data pipeline will also find this Learning Path useful. To grasp the concepts explained in this Learning Path, you must know the fundamentals of Apache Spark and Scala.

This is a reference book for Architects. This book can be helpful for those developers who wants to increase breadth of knowledge about tools and technology. If you are planning for career advancement and you are interviewing for cloud architect, this book can also be used for interview preparation purpose. You can go through this book before your interview every time, so that you will remember all the concepts before interview. As the technology is evolving very fast, new tools and technologies are coming every day. This book covers fundamental of architecting or re-architecting of the application. This book also makes you aware and provides details about tools and technology available in cloud. This book does not over explain any concepts, keeping in mind that you can complete your reading in less time. With this book, you will get lot of information in less reading time.

????

This book discusses the issues and challenges in Online Social Networks (OSNs). It highlights various aspects of OSNs consisting of novel social network strategies and the development of services using different computing models.

Moreover, the book investigates how OSNs are impacted by cutting-edge innovations.

??????????SQL Server 2005??????????,????????????????????????????????????

This book includes the outcomes of the International Conference on Advanced Intelligent Systems for Sustainable Development (AI2SD-2018), held in Tangier, Morocco on July 12–14, 2018. Presenting the latest research in the field of computing sciences and information technology, it discusses new challenges and provides valuable insights into the field, the goal being to stimulate debate, and to promote closer interaction and interdisciplinary collaboration between researchers and practitioners. Though chiefly intended for researchers and practitioners in advanced information technology management and networking, the book will also be of interest to those engaged in emerging fields such as data science and analytics, big data, internet of things, smart networked systems, artificial intelligence, expert systems and cloud computing.

[Copyright: a99781d4cb756230f09ad6dabed4b2e6](#)