

Hadoop Real World Solutions Cookbook File Type

A comprehensive guide to mastering the most advanced Hadoop 3 concepts Key Features Get to grips with the newly introduced features and capabilities of Hadoop 3 Crunch and process data using MapReduce, YARN, and a host of tools within the Hadoop ecosystem Sharpen your Hadoop skills with real-world case studies and code Book Description Apache Hadoop is one of the most popular big data solutions for distributed storage and for processing large chunks of data. With Hadoop 3, Apache promises to provide a high-performance, more fault-tolerant, and highly efficient big data processing platform, with a focus on improved scalability and increased efficiency. With this guide, you'll understand advanced concepts of the Hadoop ecosystem tool. You'll learn how Hadoop works internally, study advanced concepts of different ecosystem tools, discover solutions to real-world use cases, and understand how to secure your cluster. It will then walk you through HDFS, YARN, MapReduce, and Hadoop 3 concepts. You'll be able to address common challenges like using Kafka efficiently, designing low latency, reliable message delivery Kafka systems, and handling high data volumes. As you advance, you'll discover how to address major challenges when building an enterprise-grade messaging system, and how to use different stream processing systems along with Kafka to fulfil your enterprise goals. By the end of this book, you'll have a complete understanding of how components in the Hadoop ecosystem are effectively integrated to implement a fast and reliable data pipeline, and you'll be equipped to tackle a range of real-world problems in data pipelines. What you will learn Gain an in-depth understanding of distributed computing using Hadoop 3 Develop enterprise-grade applications using Apache Spark, Flink, and more Build

Online Library Hadoop Real World Solutions Cookbook File Type

scalable and high-performance Hadoop data pipelines with security, monitoring, and data governance Explore batch data processing patterns and how to model data in Hadoop Master best practices for enterprises using, or planning to use, Hadoop 3 as a data platform Understand security aspects of Hadoop, including authorization and authentication Who this book is for If you want to become a big data professional by mastering the advanced concepts of Hadoop, this book is for you. You'll also find this book useful if you're a Hadoop professional looking to strengthen your knowledge of the Hadoop ecosystem. Fundamental knowledge of the Java programming language and basics of Hadoop is necessary to get started with this book.

Combine the power of analytics and cloud computing for faster and efficient insights Key Features Master the concept of analytics on the cloud: and how organizations are using it Learn the design considerations and while applying a cloud analytics solution Design an end-to-end analytics pipeline on the cloud Book Description With the ongoing data explosion, more and more organizations all over the world are slowly migrating their infrastructure to the cloud. These cloud platforms also provide their distinct analytics services to help you get faster insights from your data. This book will give you an introduction to the concept of analytics on the cloud, and the different cloud services popularly used for processing and analyzing data. If you're planning to adopt the cloud analytics model for your business, this book will help you understand the design and business considerations to be kept in mind, and choose the best tools and alternatives for analytics, based on your requirements. The chapters in this book will take you through the 70+ services available in Google Cloud Platform and their implementation for practical purposes. From ingestion to processing your data, this book contains best

practices on building an end-to-end analytics pipeline on the cloud by leveraging popular concepts such as machine learning and deep learning. By the end of this book, you will have a better understanding of cloud analytics as a concept as well as a practical know-how of its implementation. What you will learn: Explore the basics of cloud analytics and the major cloud solutions. Learn how organizations are using cloud analytics to improve the ROI. Explore the design considerations while adopting cloud services. Work with the ingestion and storage tools of GCP such as Cloud Pub/Sub. Process your data with tools such as Cloud Dataproc, BigQuery, etc. Over 70 GCP tools to build an analytics engine for cloud analytics. Implement machine learning and other AI techniques on GCP. Who this book is for: This book is targeted at CIOs, CTOs, and even analytics professionals looking for various alternatives to implement their analytics pipeline on the cloud. Data professionals looking to get started with cloud-based analytics will also find this book useful. Some basic exposure to cloud platforms such as GCP will be helpful, but not mandatory.

Until now, design patterns for the MapReduce framework have been scattered among various research papers, blogs, and books. This handy guide brings together a unique collection of valuable MapReduce patterns that will save you time and effort regardless of the domain, language, or development framework you're using. Each pattern is explained in context, with pitfalls and caveats clearly identified to help you avoid common design mistakes when modeling your big data architecture. This book also provides a complete overview of MapReduce that explains its origins and implementations, and why design patterns are so important. All code examples are written for Hadoop. Summarization patterns: get a top-level view by summarizing and grouping data. Filtering patterns: view data subsets such as records

generated from one user Data organization patterns: reorganize data to work with other systems, or to make MapReduce analysis easier Join patterns: analyze different datasets together to discover interesting relationships Metapatterns: piece together several patterns to solve multi-stage problems, or to perform several analytics in the same job Input and output patterns: customize the way you use Hadoop to load or store data "A clear exposition of MapReduce programs for common data processing patterns—this book is indispensable for anyone using Hadoop." --Tom White, author of Hadoop: The Definitive Guide

Although you don't need a large computing infrastructure to process massive amounts of data with Apache Hadoop, it can still be difficult to get started. This practical guide shows you how to quickly launch data analysis projects in the cloud by using Amazon Elastic MapReduce (EMR), the hosted Hadoop framework in Amazon Web Services (AWS). Authors Kevin Schmidt and Christopher Phillips demonstrate best practices for using EMR and various AWS and Apache technologies by walking you through the construction of a sample MapReduce log analysis application. Using code samples and example configurations, you'll learn how to assemble the building blocks necessary to solve your biggest data analysis problems. Get an overview of the AWS and Apache software tools used in large-scale data analysis Go through the process of executing a Job Flow with a simple log analyzer Discover useful MapReduce patterns for filtering and analyzing data sets Use Apache Hive and Pig instead of Java to build a MapReduce Job Flow Learn the basics for using Amazon EMR to run machine learning algorithms Develop a project cost model for using Amazon EMR and other AWS tools Integrating data from multiple sources is essential in the age of big data, but it can be a challenging and time-consuming task. This handy cookbook provides dozens of ready-to-use

Online Library Hadoop Real World Solutions Cookbook File Type

recipes for using Apache Sqoop, the command-line interface application that optimizes data transfers between relational databases and Hadoop. Sqoop is both powerful and bewildering, but with this cookbook's problem-solution-discussion format, you'll quickly learn how to deploy and then apply Sqoop in your environment. The authors provide MySQL, Oracle, and PostgreSQL database examples on GitHub that you can easily adapt for SQL Server, Netezza, Teradata, or other relational systems. Transfer data from a single database table into your Hadoop ecosystem Keep table data and Hadoop in sync by importing data incrementally Import data from more than one database table Customize transferred data by calling various database functions Export generated, processed, or backed-up data from Hadoop to your database Run Sqoop within Oozie, Hadoop's specialized workflow scheduler Load data into Hadoop's data warehouse (Hive) or database (HBase) Handle installation, connection, and syntax issues common to specific database vendors

If you have interest in DynamoDB and want to know what DynamoDB is all about and become proficient in using it, this is the book for you. If you are an intermediate user who wishes to enhance your knowledge of DynamoDB, this book is aimed at you. Basic familiarity with programming, NoSQL, and cloud computing concepts would be helpful.

Designed for the way many developers work, this practical problem-solving guide balances the need for rapid development with a trusted source of information.

Over 90 hands-on recipes to help you learn and master the intricacies of Apache Hadoop 2.X, YARN, Hive, Pig, Oozie, Flume, Sqoop, Apache Spark, and Mahout About This Book Implement outstanding Machine Learning use cases on your own analytics models and processes. Solutions to common problems when working with the Hadoop ecosystem. Step-by-

step implementation of end-to-end big data use cases. Who This Book Is For Readers who have a basic knowledge of big data systems and want to advance their knowledge with hands-on recipes. What You Will Learn Installing and maintaining Hadoop 2.X cluster and its ecosystem. Write advanced Map Reduce programs and understand design patterns. Advanced Data Analysis using the Hive, Pig, and Map Reduce programs. Import and export data from various sources using Sqoop and Flume. Data storage in various file formats such as Text, Sequential, Parquet, ORC, and RC Files. Machine learning principles with libraries such as Mahout Batch and Stream data processing using Apache Spark In Detail Big data is the current requirement. Most organizations produce huge amount of data every day. With the arrival of Hadoop-like tools, it has become easier for everyone to solve big data problems with great efficiency and at minimal cost. Grasping Machine Learning techniques will help you greatly in building predictive models and using this data to make the right decisions for your organization. Hadoop Real World Solutions Cookbook gives readers insights into learning and mastering big data via recipes. The book not only clarifies most big data tools in the market but also provides best practices for using them. The book provides recipes that are based on the latest versions of Apache Hadoop 2.X, YARN, Hive, Pig, Sqoop, Flume, Apache Spark, Mahout and many more such ecosystem tools. This real-world-solution cookbook is packed with handy recipes you can apply to your own everyday issues. Each chapter provides in-depth recipes that can be referenced easily. This book provides detailed practices on the latest technologies such as YARN and Apache Spark. Readers will be able to consider themselves as big data experts on completion of this book. This guide is an invaluable tutorial if you are planning to implement a big data warehouse for your business. Style and approach An easy-to-

follow guide that walks you through world of big data. Each tool in the Hadoop ecosystem is explained in detail and the recipes are placed in such a manner that readers can implement them sequentially. Plenty of reference links are provided for advanced reading.

Analyze vast amounts of data in record time using Apache Spark with Databricks in the Cloud. Learn the fundamentals, and more, of running analytics on large clusters in Azure and AWS, using Apache Spark with Databricks on top. Discover how to squeeze the most value out of your data at a mere fraction of what classical analytics solutions cost, while at the same time getting the results you need, incrementally faster. This book explains how the confluence of these pivotal technologies gives you enormous power, and cheaply, when it comes to huge datasets. You will begin by learning how cloud infrastructure makes it possible to scale your code to large amounts of processing units, without having to pay for the machinery in advance. From there you will learn how Apache Spark, an open source framework, can enable all those CPUs for data analytics use. Finally, you will see how services such as Databricks provide the power of Apache Spark, without you having to know anything about configuring hardware or software. By removing the need for expensive experts and hardware, your resources can instead be allocated to actually finding business value in the data. This book guides you through some advanced topics such as analytics in the cloud, data lakes, data ingestion, architecture, machine learning, and tools, including Apache Spark, Apache Hadoop, Apache Hive, Python, and SQL. Valuable exercises help reinforce what you

have learned. What You Will Learn Discover the value of big data analytics that leverage the power of the cloud Get started with Databricks using SQL and Python in either Microsoft Azure or AWS Understand the underlying technology, and how the cloud and Apache Spark fit into the bigger picture See how these tools are used in the real world Run basic analytics, including machine learning, on billions of rows at a fraction of a cost or free Who This Book Is For Data engineers, data scientists, and cloud architects who want or need to run advanced analytics in the cloud. It is assumed that the reader has data experience, but perhaps minimal exposure to Apache Spark and Azure Databricks. The book is also recommended for people who want to get started in the analytics field, as it provides a strong foundation.

A Cookbook with plenty of practical recipes for different uses of Storm. If you are a Java developer with basic knowledge of real-time processing and would like to learn Storm to process unbounded streams of data in real time, then this book is for you.

Harness the power of SQL Server 2017 Integration Services to build your data integration solutions with ease About This Book Acquaint yourself with all the newly introduced features in SQL Server 2017 Integration Services Program and extend your packages to enhance their functionality This detailed, step-by-step guide covers everything you need to develop efficient data integration and data transformation solutions for your organization Who This Book Is For This book is ideal for software engineers, DW/ETL architects, and ETL developers who need to create a new, or

enhance an existing, ETL implementation with SQL Server 2017 Integration Services. This book would also be good for individuals who develop ETL solutions that use SSIS and are keen to learn the new features and capabilities in SSIS 2017. What You Will Learn Understand the key components of an ETL solution using SQL Server 2016-2017 Integration Services Design the architecture of a modern ETL solution Have a good knowledge of the new capabilities and features added to Integration Services Implement ETL solutions using Integration Services for both on-premises and Azure data Improve the performance and scalability of an ETL solution Enhance the ETL solution using a custom framework Be able to work on the ETL solution with many other developers and have common design paradigms or techniques Effectively use scripting to solve complex data issues In Detail SQL Server Integration Services is a tool that facilitates data extraction, consolidation, and loading options (ETL), SQL Server coding enhancements, data warehousing, and customizations. With the help of the recipes in this book, you'll gain complete hands-on experience of SSIS 2017 as well as the 2016 new features, design and development improvements including SCD, Tuning, and Customizations. At the start, you'll learn to install and set up SSIS as well other SQL Server resources to make optimal use of this Business Intelligence tools. We'll begin by taking you through the new features in SSIS 2016/2017 and implementing the necessary features to get a modern scalable ETL solution that fits the modern data warehouse. Through the course of chapters, you will learn how to design and build

SSIS data warehouses packages using SQL Server Data Tools. Additionally, you'll learn to develop SSIS packages designed to maintain a data warehouse using the Data Flow and other control flow tasks. You'll also be demonstrated many recipes on cleansing data and how to get the end result after applying different transformations. Some real-world scenarios that you might face are also covered and how to handle various issues that you might face when designing your packages. At the end of this book, you'll get to know all the key concepts to perform data integration and transformation. You'll have explored on-premises Big Data integration processes to create a classic data warehouse, and will know how to extend the toolbox with custom tasks and transforms. Style and approach This cookbook follows a problem-solution approach and tackles all kinds of data integration scenarios by using the capabilities of SQL Server 2016 Integration Services. This book is well supplemented with screenshots, tips, and tricks. Each recipe focuses on a particular task and is written in a very easy-to-follow manner.

Summary Hadoop in Practice, Second Edition provides over 100 tested, instantly useful techniques that will help you conquer big data, using Hadoop. This revised new edition covers changes and new features in the Hadoop core architecture, including MapReduce 2. Brand new chapters cover YARN and integrating Kafka, Impala, and Spark SQL with Hadoop. You'll also get new and updated techniques for Flume, Sqoop, and Mahout, all of which have seen major new versions recently. In short, this is the

most practical, up-to-date coverage of Hadoop available anywhere. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book It's always a good time to upgrade your Hadoop skills! Hadoop in Practice, Second Edition provides a collection of 104 tested, instantly useful techniques for analyzing real-time streams, moving data securely, machine learning, managing large-scale clusters, and taming big data using Hadoop. This completely revised edition covers changes and new features in Hadoop core, including MapReduce 2 and YARN. You'll pick up hands-on best practices for integrating Spark, Kafka, and Impala with Hadoop, and get new and updated techniques for the latest versions of Flume, Sqoop, and Mahout. In short, this is the most practical, up-to-date coverage of Hadoop available. Readers need to know a programming language like Java and have basic familiarity with Hadoop. What's Inside Thoroughly updated for Hadoop 2 How to write YARN applications Integrate real-time technologies like Storm, Impala, and Spark Predictive analytics using Mahout and RR Readers need to know a programming language like Java and have basic familiarity with Hadoop. About the Author Alex Holmes works on tough big-data problems. He is a software engineer, author, speaker, and blogger specializing in large-scale Hadoop projects. Table of Contents PART 1 BACKGROUND AND FUNDAMENTALS Hadoop in a heartbeat Introduction to YARN PART 2 DATA LOGISTICS Data serialization—working with text and beyond Organizing and optimizing data in HDFS Moving data into and out of

Hadoop PART 3 BIG DATA PATTERNS Applying MapReduce patterns to big data Utilizing data structures and algorithms at scale Tuning, debugging, and testing PART 4 BEYOND MAPREDUCE SQL on Hadoop Writing a YARN application

Learn to solve challenging data science problems by building powerful machine learning models using Python About This Book Understand which algorithms to use in a given context with the help of this exciting recipe-based guide This practical tutorial tackles real-world computing problems through a rigorous and effective approach Build state-of-the-art models and develop personalized recommendations to perform machine learning at scale Who This Book Is For This Learning Path is for Python programmers who are looking to use machine learning algorithms to create real-world applications. It is ideal for Python professionals who want to work with large and complex datasets and Python developers and analysts or data scientists who are looking to add to their existing skills by accessing some of the most powerful recent trends in data science. Experience with Python, Jupyter Notebooks, and command-line execution together with a good level of mathematical knowledge to understand the concepts is expected. Machine learning basic knowledge is also expected. What You Will Learn Use predictive modeling and apply it to real-world problems Understand how to perform market segmentation using unsupervised learning Apply your new-found skills to solve real problems, through clearly-explained code for every technique and test Compete with top data scientists by gaining a practical and theoretical

understanding of cutting-edge deep learning algorithms Increase predictive accuracy with deep learning and scalable data-handling techniques Work with modern state-of-the-art large-scale machine learning techniques Learn to use Python code to implement a range of machine learning algorithms and techniques In Detail Machine learning is increasingly spreading in the modern data-driven world. It is used extensively across many fields such as search engines, robotics, self-driving cars, and more. Machine learning is transforming the way we understand and interact with the world around us. In the first module, Python Machine Learning Cookbook, you will learn how to perform various machine learning tasks using a wide variety of machine learning algorithms to solve real-world problems and use Python to implement these algorithms. The second module, Advanced Machine Learning with Python, is designed to take you on a guided tour of the most relevant and powerful machine learning techniques and you'll acquire a broad set of powerful skills in the area of feature selection and feature engineering. The third module in this learning path, Large Scale Machine Learning with Python, dives into scalable machine learning and the three forms of scalability. It covers the most effective machine learning techniques on a map reduce framework in Hadoop and Spark in Python. This Learning Path will teach you Python machine learning for the real world. The machine learning techniques covered in this Learning Path are at the forefront of commercial practice. This Learning Path combines some of the best that Packt has to offer in one complete, curated package. It includes content from the following Packt

products: Python Machine Learning Cookbook by Prateek Joshi Advanced Machine Learning with Python by John Hearty Large Scale Machine Learning with Python by Bastiaan Sjardin, Alberto Boschetti, Luca Massaron Style and approach This course is a smooth learning path that will teach you how to get started with Python machine learning for the real world, and develop solutions to real-world problems. Through this comprehensive course, you'll learn to create the most effective machine learning techniques from scratch and more!

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

processing with the Pig query language Analyze datasets with Hive, Hadoop's data warehousing system Take advantage of HBase for structured and semi-structured data, and ZooKeeper for building distributed systems

Uses real-world bug reports (vulnerabilities in software or in this case web applications) to teach programmers and InfoSec professionals how to discover and protect vulnerabilities in web applications. Real-World Bug Hunting is a field guide to finding software bugs. Ethical hacker Peter Yaworski breaks down common types of bugs, then contextualizes them with real bug bounty reports released by hackers on companies like Twitter, Facebook, Google, Uber, and Starbucks. As you read each report, you'll gain deeper insight into how the vulnerabilities work and how you might find similar ones. Each chapter begins with an explanation of a vulnerability type, then moves into a series of real bug bounty reports that show how the bugs were found. You'll learn things like how Cross-Site Request Forgery tricks users into unknowingly submitting information to websites they are logged into; how to pass along unsafe JavaScript to execute Cross-Site Scripting; how to access another user's data via Insecure Direct Object References; how to trick websites into disclosing information with Server Side Request Forgeries; and how bugs in application logic can lead to pretty serious vulnerabilities. Yaworski also shares advice on how to write effective vulnerability reports and develop relationships with bug bounty programs, as well as recommends hacking tools that can make the job a little easier.

Online Library Hadoop Real World Solutions Cookbook File Type

A fast paced guide that will help you learn about Apache Hadoop 3 and its ecosystem
Key Features Set up, configure and get started with Hadoop to get useful insights from large data sets Work with the different components of Hadoop such as MapReduce, HDFS and YARN Learn about the new features introduced in Hadoop 3 Book
Description Apache Hadoop is a widely used distributed data platform. It enables large datasets to be efficiently processed instead of using one large computer to store and process the data. This book will get you started with the Hadoop ecosystem, and introduce you to the main technical topics, including MapReduce, YARN, and HDFS. The book begins with an overview of big data and Apache Hadoop. Then, you will set up a pseudo Hadoop development environment and a multi-node enterprise Hadoop cluster. You will see how the parallel programming paradigm, such as MapReduce, can solve many complex data processing problems. The book also covers the important aspects of the big data software development lifecycle, including quality assurance and control, performance, administration, and monitoring. You will then learn about the Hadoop ecosystem, and tools such as Kafka, Sqoop, Flume, Pig, Hive, and HBase. Finally, you will look at advanced topics, including real time streaming using Apache Storm, and data analytics using Apache Spark. By the end of the book, you will be well versed with different configurations of the Hadoop 3 cluster. What you will learn Store and analyze data at scale using HDFS, MapReduce and YARN Install and configure Hadoop 3 in different modes Use Yarn effectively to run different applications on

Hadoop based platform Understand and monitor how Hadoop cluster is managed Consume streaming data using Storm, and then analyze it using Spark Explore Apache Hadoop ecosystem components, such as Flume, Sqoop, HBase, Hive, and Kafka Who this book is for Aspiring Big Data professionals who want to learn the essentials of Hadoop 3 will find this book to be useful. Existing Hadoop users who want to get up to speed with the new features introduced in Hadoop 3 will also benefit from this book. Having knowledge of Java programming will be an added advantage.

If you are ready to dive into the MapReduce framework for processing large datasets, this practical book takes you step by step through the algorithms and tools you need to build distributed MapReduce applications with Apache Hadoop or Apache Spark. Each chapter provides a recipe for solving a massive computational problem, such as building a recommendation system. You'll learn how to implement the appropriate MapReduce solution with code that you can use in your projects. Dr. Mahmoud Parsian covers basic design patterns, optimization techniques, and data mining and machine learning solutions for problems in bioinformatics, genomics, statistics, and social network analysis. This book also includes an overview of MapReduce, Hadoop, and Spark. Topics include: Market basket analysis for a large set of transactions Data mining algorithms (K-means, KNN, and Naive Bayes) Using huge genomic data to

sequence DNA and RNA Naive Bayes theorem and Markov chains for data and market prediction Recommendation algorithms and pairwise document similarity Linear regression, Cox regression, and Pearson correlation Allelic frequency and mining DNA Social network analysis (recommendation systems, counting triangles, sentiment analysis)

Describes the features and functions of Apache Hive, the data infrastructure for Hadoop.

If you are a system or application developer interested in learning how to solve practical problems using the Hadoop framework, then this book is ideal for you. You are expected to be familiar with the Unix/Linux command-line interface and have some experience with the Java programming language. Familiarity with Hadoop would be a plus.

Solve all your Spring 5 problems using complete and real-world code examples. When you start a new project, you'll be able to copy the code and configuration files from this book, and then modify them for your needs. This can save you a great deal of work over creating a project from scratch. The recipes in Spring 5 Recipes cover Spring fundamentals such as Spring IoC container, Spring AOP/AspectJ, and more. Other recipes include Spring enterprise solutions for topics such as Spring Java EE integration, Spring Integration, Spring Batch, Spring

Remoting, messaging, transactions, and working with big data and the cloud using Hadoop and MongoDB. Finally, Spring web recipes cover Spring MVC, other dynamic scripting, integration with the popular Grails Framework (and Groovy), REST/web services, and more. You'll also see recipes on new topics such as Spring Framework 5, reactive Spring, Spring 5 microservices, the functional web framework and much more. This book builds upon the best-selling success of the previous editions and focuses on the latest Spring Framework features for building enterprise Java applications. What You'll Learn Get reusable code recipes and snippets for core Spring, annotations and other development tools Access Spring MVC for web development Work with Spring REST and microservices for web services development and integration into your enterprise Java applications Use Spring Batch, NoSQL and big data for building and integrating various cloud computing services and resources Integrate Java Enterprise Edition and other Java APIs for use in Spring Use Grails code and much more Who This Book Is For Experienced Java and Spring programmers. Get command of your organizational Big Data using the power of data science and analytics Key Features A perfect companion to boost your Big Data storing, processing, analyzing skills to help you take informed business decisions Work with the best tools such as Apache Hadoop, R, Python, and Spark for NoSQL

platforms to perform massive online analyses Get expert tips on statistical inference, machine learning, mathematical modeling, and data visualization for Big Data Book Description Big Data analytics relates to the strategies used by organizations to collect, organize and analyze large amounts of data to uncover valuable business insights that otherwise cannot be analyzed through traditional systems. Crafting an enterprise-scale cost-efficient Big Data and machine learning solution to uncover insights and value from your organization's data is a challenge. Today, with hundreds of new Big Data systems, machine learning packages and BI Tools, selecting the right combination of technologies is an even greater challenge. This book will help you do that. With the help of this guide, you will be able to bridge the gap between the theoretical world of technology with the practical ground reality of building corporate Big Data and data science platforms. You will get hands-on exposure to Hadoop and Spark, build machine learning dashboards using R and R Shiny, create web-based apps using NoSQL databases such as MongoDB and even learn how to write R code for neural networks. By the end of the book, you will have a very clear and concrete understanding of what Big Data analytics means, how it drives revenues for organizations, and how you can develop your own Big Data analytics solution using different tools and methods articulated in this book. What you will learn -

Get a 360-degree view into the world of Big Data, data science and machine learning - Broad range of technical and business Big Data analytics topics that caters to the interests of the technical experts as well as corporate IT executives - Get hands-on experience with industry-standard Big Data and machine learning tools such as Hadoop, Spark, MongoDB, KDB+ and R - Create production-grade machine learning BI Dashboards using R and R Shiny with step-by-step instructions - Learn how to combine open-source Big Data, machine learning and BI Tools to create low-cost business analytics applications - Understand corporate strategies for successful Big Data and data science projects - Go beyond general-purpose analytics to develop cutting-edge Big Data applications using emerging technologies Who this book is for The book is intended for existing and aspiring Big Data professionals who wish to become the go-to person in their organization when it comes to Big Data architecture, analytics, and governance. While no prior knowledge of Big Data or related technologies is assumed, it will be helpful to have some programming experience.

The professional's one-stop guide to this open-source, Java-based big data framework Professional Hadoop is the complete reference and resource for experienced developers looking to employ Apache Hadoop in real-world settings. Written by an expert team of certified Hadoop developers, committers, and

Summit speakers, this book details every key aspect of Hadoop technology to enable optimal processing of large data sets. Designed expressly for the professional developer, this book skips over the basics of database development to get you acquainted with the framework's processes and capabilities right away. The discussion covers each key Hadoop component individually, culminating in a sample application that brings all of the pieces together to illustrate the cooperation and interplay that make Hadoop a major big data solution. Coverage includes everything from storage and security to computing and user experience, with expert guidance on integrating other software and more. Hadoop is quickly reaching significant market usage, and more and more developers are being called upon to develop big data solutions using the Hadoop framework. This book covers the process from beginning to end, providing a crash course for professionals needing to learn and apply Hadoop quickly. Configure storage, UE, and in-memory computing Integrate Hadoop with other programs including Kafka and Storm Master the fundamentals of Apache Big Top and Ignite Build robust data security with expert tips and advice Hadoop's popularity is largely due to its accessibility. Open-source and written in Java, the framework offers almost no barrier to entry for experienced database developers already familiar with the skills and requirements real-world programming entails. Professional Hadoop

gives you the practical information and framework-specific skills you need quickly.

Getting Started With Amazon Redshift is a step-by-step, practical guide to the world of Redshift. Learn to load, manage, and query data on Redshift. This book is for CIOs, enterprise architects, developers, and anyone else who needs to get familiar with RedShift. The CIO will gain an understanding of what their technical staff is working on; the technical implementation personnel will get an in-depth view of the technology, and what it will take to implement their own solutions. The go-to guidebook for deploying Big Data solutions with Hadoop Today's enterprise architects need to understand how the Hadoop frameworks and APIs fit together, and how they can be integrated to deliver real-world solutions. This book is a practical, detailed guide to building and implementing those solutions, with code-level instruction in the popular Wrox tradition. It covers storing data with HDFS and Hbase, processing data with MapReduce, and automating data processing with Oozie. Hadoop security, running Hadoop with Amazon Web Services, best practices, and automating Hadoop processes in real time are also covered in depth. With in-depth code examples in Java and XML and the latest on recent additions to the Hadoop ecosystem, this complete resource also covers the use of APIs, exposing their inner workings and allowing architects and

developers to better leverage and customize them. The ultimate guide for developers, designers, and architects who need to build and deploy Hadoop applications. Covers storing and processing data with various technologies, automating data processing, Hadoop security, and delivering real-time solutions. Includes detailed, real-world examples and code-level guidelines. Explains when, why, and how to use these tools effectively. Written by a team of Hadoop experts in the programmer-to-programmer Wrox style. Professional Hadoop Solutions is the reference enterprise architects and developers need to maximize the power of Hadoop.

This book will guide you from the basic functionality of Tableau 2019.x to complex deployment. It is full of useful recipes from industry experts, who will help you master your Tableau skills. The complexity of tasks increases gradually, all the way to mastering advanced functionality through bite-sized, detailed recipes.

Realistic, simple code examples to solve problems at scale with Hadoop and related technologies.

Easy, hands-on recipes to help you understand Hive and its integration with frameworks that are used widely in today's big data world. About This Book: Grasp a complete reference of different Hive topics. Get to know the latest recipes in development in Hive including CRUD operations. Understand Hive internals and

integration of Hive with different frameworks used in today's world. Who This Book Is For The book is intended for those who want to start in Hive or who have basic understanding of Hive framework. Prior knowledge of basic SQL command is also required What You Will Learn Learn different features and offering on the latest Hive Understand the working and structure of the Hive internals Get an insight on the latest development in Hive framework Grasp the concepts of Hive Data Model Master the key concepts like Partition, Buckets and Statistics Know how to integrate Hive with other frameworks such as Spark, Accumulo, etc In Detail Hive was developed by Facebook and later open sourced in Apache community. Hive provides SQL like interface to run queries on Big Data frameworks. Hive provides SQL like syntax also called as HiveQL that includes all SQL capabilities like analytical functions which are the need of the hour in today's Big Data world. This book provides you easy installation steps with different types of metastores supported by Hive. This book has simple and easy to learn recipes for configuring Hive clients and services. You would also learn different Hive optimizations including Partitions and Bucketing. The book also covers the source code explanation of latest Hive version. Hive Query Language is being used by other frameworks including spark. Towards the end you will cover integration of Hive with these frameworks. Style and approach Starting with the basics and covering the core concepts with the practical usage, this book is a complete guide to learn and explore Hive offerings.

Over 90 hands-on recipes to design Internet scalable web and mobile applications with Amazon DynamoDB About This Book Construct top-notch mobile and web applications with the Internet scalable NoSQL database and host it on cloud Integrate your applications with other AWS services like AWS EMR, AWS S3, AWS Redshift, and AWS CloudSearch etc. in order to achieve a one-stop application stack Step-by-step implementation guide that provides real-world use with hands-on recipes Who This Book Is For This book is intended for those who have a basic understanding of AWS services and want to take their knowledge to the next level by getting their hands dirty with coding recipes in DynamoDB. What You Will Learn Design DynamoDB tables to achieve high read and write throughput Discover best practices like caching, exponential back-offs and auto-retries, storing large items in AWS S3, storing compressed data etc. Effectively use DynamoDB Local in order to make your development smooth and cost effective Implement cost effective best practices to reduce the burden of DynamoDB charges Create and maintain secondary indexes to support improved data access Integrate various other AWS services like AWS EMR, AWS CloudSearch, AWS Pipeline etc. with DynamoDB In Detail AWS DynamoDB is an excellent example of a production-ready NoSQL database. In recent years, DynamoDB has been able to attract many customers because of its features like high-availability, reliability and infinite scalability. DynamoDB can be easily integrated with massive data crunching tools like Hadoop /EMR, which is an essential part of this data-driven world

and hence it is widely accepted. The cost and time-efficient design makes DynamoDB stand out amongst its peers. The design of DynamoDB is so neat and clean that it has inspired many NoSQL databases to simply follow it. This book will get your hands on some engineering best practices DynamoDB engineers use, which can be used in your day-to-day life to build robust and scalable applications. You will start by operating with DynamoDB tables and learn to manipulate items and manage indexes. You will also discover how to easily integrate applications with other AWS services like EMR, S3, CloudSearch, RedShift etc. A couple of chapters talk in detail about how to use DynamoDB as a backend database and hosting it on AWS ElasticBean. This book will also focus on security measures of DynamoDB as well by providing techniques on data encryption, masking etc. By the end of the book you'll be adroit in designing web and mobile applications using DynamoDB and host it on cloud. Style and approach An easy-to-follow guide, full of real-world examples, which takes you through the world of DynamoDB following a step-by-step, problem-solution based approach.

Hadoop in Action teaches readers how to use Hadoop and write MapReduce programs. The intended readers are programmers, architects, and project managers who have to process large amounts of data offline. Hadoop in Action will lead the reader from obtaining a copy of Hadoop to setting it up in a cluster and writing data analytic programs. The book begins by making the basic idea of Hadoop and MapReduce easier to grasp by applying the default Hadoop installation to a few easy-to-follow tasks,

such as analyzing changes in word frequency across a body of documents. The book continues through the basic concepts of MapReduce applications developed using Hadoop, including a close look at framework components, use of Hadoop for a variety of data analysis tasks, and numerous examples of Hadoop in action. Hadoop in Action will explain how to use Hadoop and present design patterns and practices of programming MapReduce. MapReduce is a complex idea both conceptually and in its implementation, and Hadoop users are challenged to learn all the knobs and levers for running Hadoop. This book takes you beyond the mechanics of running Hadoop, teaching you to write meaningful programs in a MapReduce framework. This book assumes the reader will have a basic familiarity with Java, as most code examples will be written in Java. Familiarity with basic statistical concepts (e.g. histogram, correlation) will help the reader appreciate the more advanced data processing examples. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.

Big Data Analytics with R and Hadoop is a tutorial style book that focuses on all the powerful big data tasks that can be achieved by integrating R and Hadoop. This book is ideal for R developers who are looking for a way to perform big data analytics with Hadoop. This book is also aimed at those who know Hadoop and want to build some intelligent applications over Big data with R packages. It would be helpful if readers have basic knowledge of R.

While many companies ponder implementation details such as distributed processing engines and algorithms for data analysis, this practical book takes a much wider view of big data development, starting with initial planning and moving diligently toward execution. Authors Ted Malaska and Jonathan Seidman guide you through the major components necessary to start, architect, and develop successful big data projects. Everyone from CIOs and COOs to lead architects and developers will explore a variety of big data architectures and applications, from massive data pipelines to web-scale applications. Each chapter addresses a piece of the software development life cycle and identifies patterns to maximize long-term success throughout the life of your project. Start the planning process by considering the key data project types Use guidelines to evaluate and select data management solutions Reduce risk related to technology, your team, and vague requirements Explore system interface design using APIs, REST, and pub/sub systems Choose the right distributed storage system for your big data system Plan and implement metadata collections for your data architecture Use data pipelines to ensure data integrity from source to final storage Evaluate the attributes of various engines for processing the data you collect

This book is intended for developers who have some familiarity with Apache Karaf and who want a quick reference for practical, proven tips on how to perform common tasks such as configuring Pax modules deployed in Apache Karaf, Extending `HttpService` with Apache Karaf. You should have working knowledge of Apache karaf, as the book

provides a deeper understanding of the capabilities of Apache Karaf.

Big data is a relative term describing a situation where the volume, velocity and variety of data exceed an organization's storage or compute capacity for accurate and timely decision making . Big data is not a single technology but a combination of old and new technologies that helps companies gain actionable insight. Therefore, big data is the capability to manage a huge volume of disparate data, at the right speed, and within the right time frame to allow real-time analysis and reaction. As we note earlier in this chapter, big data is typically broken down by three characteristics: Volume: How much data Velocity: How fast that data is processed Variety: The various types of data Although it's convenient to simplify big data into the three Vs, it can be misleading and overly simplistic. For example, you may be managing a relatively small amount of very disparate, complex data or you may be processing a huge volume of very simple data. That simple data may be all structured or all unstructured. Even more important is the fourth V: veracity. How accurate is that data in predicting business value? Do the results of a big data analysis actually make sense? Determining relevant data is key to delivering value from massive amounts of data. However, big data is defined less by volume - which is a constantly moving target - than by its ever-increasing variety, velocity, variability and complexity .

If you've been asked to maintain large and complex Hadoop clusters, this book is a must. Demand for operations-specific material has skyrocketed now that Hadoop is

becoming the de facto standard for truly large-scale data processing in the data center. Eric Sammer, Principal Solution Architect at Cloudera, shows you the particulars of running Hadoop in production, from planning, installing, and configuring the system to providing ongoing maintenance. Rather than run through all possible scenarios, this pragmatic operations guide calls out what works, as demonstrated in critical deployments. Get a high-level overview of HDFS and MapReduce: why they exist and how they work Plan a Hadoop deployment, from hardware and OS selection to network requirements Learn setup and configuration details with a list of critical properties Manage resources by sharing a cluster across multiple groups Get a runbook of the most common cluster maintenance tasks Monitor Hadoop clusters—and learn troubleshooting with the help of real-world war stories Use basic tools and techniques to handle backup and catastrophic failure

If you are a Big Data enthusiast and wish to use Hadoop v2 to solve your problems, then this book is for you. This book is for Java programmers with little to moderate knowledge of Hadoop MapReduce. This is also a one-stop reference for developers and system admins who want to quickly get up to speed with using Hadoop v2. It would be helpful to have a basic knowledge of software development using Java and a basic working knowledge of Linux.

Dive into the world of SQL on Hadoop and get the most out of your Hive data warehouses. This book is your go-to resource for using Hive: authors Scott Shaw,

Ankur Gupta, David Kjerrumgaard, and Andreas Francois Vermeulen take you through learning HiveQL, the SQL-like language specific to Hive, to analyze, export, and massage the data stored across your Hadoop environment. From deploying Hive on your hardware or virtual machine and setting up its initial configuration to learning how Hive interacts with Hadoop, MapReduce, Tez and other big data technologies, Practical Hive gives you a detailed treatment of the software. In addition, this book discusses the value of open source software, Hive performance tuning, and how to leverage semi-structured and unstructured data. What You Will Learn Install and configure Hive for new and existing datasets Perform DDL operations Execute efficient DML operations Use tables, partitions, buckets, and user-defined functions Discover performance tuning tips and Hive best practices Who This Book Is For Developers, companies, and professionals who deal with large amounts of data and could use software that can efficiently manage large volumes of input. It is assumed that readers have the ability to work with SQL.

Definitive guide to lightning fast data processing for distributed systems with Apache Flink About This Book* Build your expertise in processing realtime data with Apache Flink and its ecosystem* Gain insights into the working of all components of Apache Flink such as FlinkML, Gelly, and Table API Filled with real world use cases,* Your guide to take advantage of Apache Flink for solving real world problems Who This Book Is For Big data developers who are looking to process batch and real-time data on

distributed systems. Basic knowledge of Hadoop and big data is assumed. Reasonable knowledge of Java or Scala is expected.

What You Will Learn*

- Learn how to build end to end real time analytics projects*
- Integrate with existing big data stack and utilize existing infrastructure.*
- Build predictive analytics applications using FlinkML*
- Use graph library to perform graph querying and search.

In Detail

With the advent of massive computer systems, organizations in different domains generate large amounts of data at a realtime basis. The latest entrant to big data processing, Apache Flink, is designed to process continuous streams of data at a lightning fast pace. This book will be your definitive guide to batch and stream data processing with Apache Flink. The book begins with introducing the Apache Flink ecosystem, setting it up and using the DataSet and DataStream API for processing batch and streaming datasets. Bringing the power of SQL to Flink, this book will then explore the Table API for querying and manipulating data. In the latter half of the book, readers will get to learn the remaining ecosystem of Apache Flink to achieve complex tasks such as event processing, machine learning, and graph processing. The final part of the book would consist of topics such as scaling Flink solutions, performance optimization and integrating Flink with other tools such as ElasticSearch.

Whether you want to dive deeper into Apache Flink, or want to investigate how to get more out of this powerful technology, you'll find everything inside

Utilize R to uncover hidden patterns in your Big Data About This Book

Perform computational analyses on Big Data to generate meaningful results Get a practical

knowledge of R programming language while working on Big Data platforms like Hadoop, Spark, H2O and SQL/NoSQL databases, Explore fast, streaming, and scalable data analysis with the most cutting-edge technologies in the market Who This Book Is For This book is intended for Data Analysts, Scientists, Data Engineers, Statisticians, Researchers, who want to integrate R with their current or future Big Data workflows. It is assumed that readers have some experience in data analysis and understanding of data management and algorithmic processing of large quantities of data, however they may lack specific skills related to R. What You Will Learn Learn about current state of Big Data processing using R programming language and its powerful statistical capabilities Deploy Big Data analytics platforms with selected Big Data tools supported by R in a cost-effective and time-saving manner Apply the R language to real-world Big Data problems on a multi-node Hadoop cluster, e.g. electricity consumption across various socio-demographic indicators and bike share scheme usage Explore the compatibility of R with Hadoop, Spark, SQL and NoSQL databases, and H2O platform In Detail Big Data analytics is the process of examining large and complex data sets that often exceed the computational capabilities. R is a leading programming language of data science, consisting of powerful functions to tackle all problems related to Big Data processing. The book will begin with a brief introduction to the Big Data world and its current industry standards. With introduction to the R language and presenting its development, structure, applications in real world,

and its shortcomings. Book will progress towards revision of major R functions for data management and transformations. Readers will be introduced to Cloud based Big Data solutions (e.g. Amazon EC2 instances and Amazon RDS, Microsoft Azure and its HDInsight clusters) and also provide guidance on R connectivity with relational and non-relational databases such as MongoDB and HBase etc. It will further expand to include Big Data tools such as Apache Hadoop ecosystem, HDFS and MapReduce frameworks. Also other R compatible tools such as Apache Spark, its machine learning library Spark MLlib, as well as H2O. Style and approach This book will serve as a practical guide to tackling Big Data problems using R programming language and its statistical environment. Each section of the book will present you with concise and easy-to-follow steps on how to process, transform and analyse large data sets.

Data is arriving faster than you can process it and the overall volumes keep growing at a rate that keeps you awake at night. Hadoop can help you tame the data beast. Effective use of Hadoop however requires a mixture of programming, design, and system administration skills. "Hadoop Beginner's Guide" removes the mystery from Hadoop, presenting Hadoop and related technologies with a focus on building working systems and getting the job done, using cloud services to do so when it makes sense. From basic concepts and initial setup through developing applications and keeping the system running as the data grows, the book gives the understanding needed to effectively use Hadoop to solve real world problems. Starting with the basics of

Online Library Hadoop Real World Solutions Cookbook File Type

installing and configuring Hadoop, the book explains how to develop applications, maintain the system, and how to use additional products to integrate with other systems. While learning different ways to develop applications to run on Hadoop the book also covers tools such as Hive, Sqoop, and Flume that show how Hadoop can be integrated with relational databases and log collection. In addition to examples on Hadoop clusters on Ubuntu uses of cloud services such as Amazon, EC2 and Elastic MapReduce are covered.

Over 90 hands-on recipes to help you learn and master the intricacies of Apache Hadoop 2.X, YARN, Hive, Pig, Oozie, Flume, Sqoop, Apache Spark, and Mahout

About This Book- Implement outstanding Machine Learning use cases on your own analytics models and processes.- Solutions to common problems when working with the Hadoop ecosystem.- Step-by-step implementation of end-to-end big data use cases.

Who This Book Is For Readers who have a basic knowledge of big data systems and want to advance their knowledge with hands-on recipes.

What You Will Learn- Installing and maintaining Hadoop 2.X cluster and its ecosystem.- Write advanced Map Reduce programs and understand design patterns.- Advanced Data Analysis using the Hive, Pig, and Map Reduce programs.- Import and export data from various sources using Sqoop and Flume.- Data storage in various file formats such as Text, Sequential, Parquet, ORC, and RC Files.- Machine learning principles with libraries such as Mahout- Batch and Stream data processing using Apache Spark

In Detail Big data is the

current requirement. Most organizations produce huge amount of data every day. With the arrival of Hadoop-like tools, it has become easier for everyone to solve big data problems with great efficiency and at minimal cost. Grasping Machine Learning techniques will help you greatly in building predictive models and using this data to make the right decisions for your organization. Hadoop Real World Solutions Cookbook gives readers insights into learning and mastering big data via recipes. The book not only clarifies most big data tools in the market but also provides best practices for using them. The book provides recipes that are based on the latest versions of Apache Hadoop 2.X, YARN, Hive, Pig, Sqoop, Flume, Apache Spark, Mahout and many more such ecosystem tools. This real-world-solution cookbook is packed with handy recipes you can apply to your own everyday issues. Each chapter provides in-depth recipes that can be referenced easily. This book provides detailed practices on the latest technologies such as YARN and Apache Spark. Readers will be able to consider themselves as big data experts on completion of this book. This guide is an invaluable tutorial if you are planning to implement a big data warehouse for your business. Style and approach An easy-to-follow guide that walks you through world of big data. Each tool in the Hadoop ecosystem is explained in detail and the recipes are placed in such a manner that readers can implement them sequentially. Plenty of reference links are provided for advanced reading.

[Copyright: 162b3f17fa3740768ecedd5fb51b478f](#)