

## Hadoop Mapreduce V2 Cookbook Second Edition

Unlock the power of your data with Hadoop 2.X ecosystem and its data warehousing techniques across large data sets About This Book Conquer the mountain of data using Hadoop 2.X tools The authors succeed in creating a context for Hadoop and its ecosystem Hadoop on examples and recipes giving the bigger picture and helping you to master Hadoop 2.X data processing platforms Overcome the challenging data processing problems using this exhaustive course with Hadoop 2.X Who This Book Is For This course is for Java developers, who know scripting, wanting a career shift to Hadoop - Big Data segment of the IT industry. So if you are a novice in Hadoop or an expert, this book will make you reach the most advanced level in Hadoop. You will learn Best practices for setup and configuration of Hadoop clusters, tailoring the system to the problem at hand. Integration with relational databases, using Hive for SQL queries and Sqoop for data transfer. Installing Hadoop 2.X cluster and its ecosystem. Advanced Data Analysis using the Hive, Pig, and Map Reduce programs. Machine learning principles with libraries such as Mahout and Batch and Stream data processing using Apache Spark. Understand the changes involved in the process in the move from Hadoop 1.0 to Hadoop 2.0. Dive into YARN and Storm and use YARN to integrate Storm with Hadoop. Deploy Hadoop on Amazon Elastic MapReduce and Discard HDFS replacements and learn about HDFS Federation. In Detail As Marc Andreessen has said "Data is eating the world", which can be witnessed today being the age of Big Data, businesses are producing data in huge volumes every day and this rise in tide of data need to be organized and analyzed in a more secured way. With proper and effective use of Hadoop, you can build new-improved models, and based on that you will be able to make the right decisions. The first module, Hadoop beginners Guide will walk you through an understanding Hadoop with very detailed instructions and how to go about using it. Commands are explained using sections called "What just happened" for more clarity and understanding. The second module, Hadoop Real World Solutions Cookbook, 2nd edition, is an essential tutorial to effectively implement a big data warehouse in your business, where you get detailed practices on the latest technologies such as YARN and Spark. Big data has become a key basis of competition and the new waves of productivity growth. Hence, once you get familiar with the basics and implement the end-to-end big data use cases, you will start exploring the third module, Mastering Hadoop. So, now the question is if you need to broaden your Hadoop skill set to the next level after you nail the basics and the advance concepts, then this course is indispensable. When you finish this course, you will be able to tackle the real-world scenarios and become a big data expert using the tools and techniques available in the Hadoop ecosystem. Style and approach This course covers everything right from the basic concepts of Hadoop till you master the advance mechanisms to become a big data expert. The goal here is to help you learn the basic essentials using the step-by-step tutorials and from there moving toward the recipes with various real-world solutions for you. It covers all the important aspects of Hadoop from system designing and configuring Hadoop, machine learning principles with various libraries with chapters illustrated with code fragments and schematic diagrams. This is a compendious course to explore Hadoop from the basics to the most advanced techniques available in Hadoop 2.X.

This book is intended for software developers, system architects and analysts, big data project managers, and data scientists who wish to deploy big data solutions using the Cascading framework. You must have a basic understanding of the big data paradigm and should be familiar with Java development techniques. 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 book is a must-read for system 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.

Practical Graph Analytics with Apache Giraph helps you build data mining and machine learning applications using the Apache Foundation's Giraph framework for graph processing. This is the same framework as used by Facebook, Google, and other social media analytics operations to derive business value from vast amounts of interconnected data points. Graphs arise in a wealth of data scenarios and describe the connections that are naturally formed in both digital and real worlds. Examples of such connections abound in online social networks such as Facebook and Twitter, among users who rate movies from services like Netflix and Amazon Prime, and are useful even in the context of biological networks for scientific research. Whether in the context of business or science, viewing data as connected adds value by increasing the amount of information available to be drawn from that data and put to use in generating new revenue or scientific opportunities. Apache Giraph offers a simple yet flexible programming model targeted to graph algorithms and designed to scale easily to accommodate massive amounts of data. Originally developed at Yahoo!, Giraph is now a top top-level project at the Apache Foundation, and it enlists contributors from companies such as Facebook, LinkedIn, and Twitter. Practical Graph Analytics with Apache Giraph brings the power of Apache Giraph to you, showing how to harness the power of graph processing for your own data by building sophisticated graph analytics applications using the very same framework that is relied upon by some of the largest players in the industry today.

This book highlights recent research on intelligent systems and nature-inspired computing. It presents 62 selected papers from the 19th International Conference on Intelligent Systems Design and Applications (ISDA 2019), which was held online. The ISDA is a premier conference in the field of computational intelligence, and the latest installment brought together researchers, engineers and practitioners whose work involves intelligent systems and their applications in industry. Including contributions by authors from 33 countries, the book offers a valuable reference guide for all researchers, students and practitioners in the fields of Computer Science and Engineering.

**Professional Hadoop Solutions**  
**Programming Hive**  
**Hadoop Data Pressing and Modelling**  
**Hadoop Real-World Solutions Cookbook**  
**Hadoop Application Architectures**  
*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 exist they and how they work. Best practices for setting up and configuring Hadoop clusters, tailoring the system to the problem at hand. Integration with relational databases, using Hive for SQL queries and Sqoop for data transfer. Installing Hadoop 2.X cluster and its ecosystem. Advanced Data Analysis using the Hive, Pig, and Map Reduce programs. Machine learning principles with libraries such as Mahout and Batch and Stream data processing using Apache Spark. Understand the changes involved in the process in the move from Hadoop 1.0 to Hadoop 2.0. Dive into YARN and Storm and use YARN to integrate Storm with Hadoop. Deploy Hadoop on Amazon Elastic MapReduce and Discard HDFS replacements and learn about HDFS Federation. In Detail As Marc Andreessen has said "Data is eating the world", which can be witnessed today being the age of Big Data, businesses are producing data in huge volumes every day and this rise in tide of data need to be organized and analyzed in a more secured way. With proper and effective use of Hadoop, you can build new-improved models, and based on that you will be able to make the right decisions. The first module, Hadoop beginners Guide will walk you through an understanding Hadoop with very detailed instructions and how to go about using it. Commands are explained using sections called "What just happened" for more clarity and understanding. The second module, Hadoop Real World Solutions Cookbook, 2nd edition, is an essential tutorial to effectively implement a big data warehouse in your business, where you get detailed practices on the latest technologies such as YARN and Spark. Big data has become a key basis of competition and the new waves of productivity growth. Hence, once you get familiar with the basics and implement the end-to-end big data use cases, you will start exploring the third module, Mastering Hadoop. So, now the question is if you need to broaden your Hadoop skill set to the next level after you nail the basics and the advance concepts, then this course is indispensable. When you finish this course, you will be able to tackle the real-world scenarios and become a big data expert using the tools and techniques available in the Hadoop ecosystem. Style and approach This course covers everything right from the basic concepts of Hadoop till you master the advance mechanisms to become a big data expert. The goal here is to help you learn the basic essentials using the step-by-step tutorials and from there moving toward the recipes with various real-world solutions for you. It covers all the important aspects of Hadoop from system designing and configuring Hadoop, machine learning principles with various libraries with chapters illustrated with code fragments and schematic diagrams. This is a compendious course to explore Hadoop from the basics to the most advanced techniques available in Hadoop 2.X.*

**Professional Hadoop Solutions**  
**Programming Hive**  
**Hadoop Data Pressing and Modelling**  
**Hadoop Real-World Solutions Cookbook**  
**Hadoop Application Architectures**  
*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 exist they and how they work. Best practices for setting up and configuring Hadoop clusters, tailoring the system to the problem at hand. Integration with relational databases, using Hive for SQL queries and Sqoop for data transfer. Installing Hadoop 2.X cluster and its ecosystem. Advanced Data Analysis using the Hive, Pig, and Map Reduce programs. Machine learning principles with libraries such as Mahout and Batch and Stream data processing using Apache Spark. Understand the changes involved in the process in the move from Hadoop 1.0 to Hadoop 2.0. Dive into YARN and Storm and use YARN to integrate Storm with Hadoop. Deploy Hadoop on Amazon Elastic MapReduce and Discard HDFS replacements and learn about HDFS Federation. In Detail As Marc Andreessen has said "Data is eating the world", which can be witnessed today being the age of Big Data, businesses are producing data in huge volumes every day and this rise in tide of data need to be organized and analyzed in a more secured way. With proper and effective use of Hadoop, you can build new-improved models, and based on that you will be able to make the right decisions. The first module, Hadoop beginners Guide will walk you through an understanding Hadoop with very detailed instructions and how to go about using it. Commands are explained using sections called "What just happened" for more clarity and understanding. The second module, Hadoop Real World Solutions Cookbook, 2nd edition, is an essential tutorial to effectively implement a big data warehouse in your business, where you get detailed practices on the latest technologies such as YARN and Spark. Big data has become a key basis of competition and the new waves of productivity growth. Hence, once you get familiar with the basics and implement the end-to-end big data use cases, you will start exploring the third module, Mastering Hadoop. So, now the question is if you need to broaden your Hadoop skill set to the next level after you nail the basics and the advance concepts, then this course is indispensable. When you finish this course, you will be able to tackle the real-world scenarios and become a big data expert using the tools and techniques available in the Hadoop ecosystem. Style and approach This course covers everything right from the basic concepts of Hadoop till you master the advance mechanisms to become a big data expert. The goal here is to help you learn the basic essentials using the step-by-step tutorials and from there moving toward the recipes with various real-world solutions for you. It covers all the important aspects of Hadoop from system designing and configuring Hadoop, machine learning principles with various libraries with chapters illustrated with code fragments and schematic diagrams. This is a compendious course to explore Hadoop from the basics to the most advanced techniques available in Hadoop 2.X.*

**Apache Spark 2.X Cookbook**  
**Big Data with Hadoop MapReduce**  
**A Guide for Developers and Administrators**  
**Moving Beyond MapReduce and Batch Processing with Apache Hadoop 2**  
If you want to discover one of the latest tools designed to produce stunning Big Data insights, this book features everything you need to get to grips with your data. Whether you are a data architect, developer, or a business strategist, HDInsight adds value in everything from development, administration, and reporting.

Hadoop has changed the way large data sets are analyzed, stored, transferred, and processed. At such low cost, it provides benefits like supports partial failure, fault tolerance, consistency, scalability, flexible schema, and so on. It also supports cloud computing. More and more number of individuals are looking forward to mastering their Hadoop skills. While initiating with Hadoop, most users are unsure about how to proceed with Hadoop. They are not aware of what are the pre-requisites or data structure they should be familiar with. Or How to make the most efficient use of Hadoop and its ecosystem. To help them with all these queries and other issues this e-book is designed. The book gives insights into many of Hadoop libraries and packages that are not known to many Big Data analysts and Architects. The e-book also tells you about Hadoop MapReduce and HDFS. The example in the e-book is well chosen and demonstrates how to control Hadoop ecosystem through various shell commands. With this book, users will gain expertise in Hadoop technology and its related components. The book also tells you with the basic concepts and Reason through this book, you will also acquire knowledge on the Hadoop Certifications like COAH and COOH. It is a definitive guide to Hadoop. Table Of Contents: Chapter 1: What is Hadoop? Chapter 2: Categories Of Big Data 2. Characteristics Of Big Data 3. Advantages Of Big Data Processing Chapter 2: Introduction to Hadoop 1. Components of Hadoop 2. Features Of Hadoop 3. Network Topology In Hadoop Chapter 3: Hadoop Installation Chapter 4: HDFS 1. Read Operation 2. Write Operation 3. Access HDFS using JAVA API 4. Access HDFS Using COMMAND-LINE INTERFACE Chapter 5: MapReduce 1. How MapReduce works 2. How MapReduce Organizes Work? Chapter 6: First Program 1. Understanding MapReducer Code 2. Explanation of SalesMapReducer Class 3. Explanation of SalesCountryReducer Class 4. Explanation of SalesCountryDriver Class Chapter 7: Counters & Joins In MapReduce 1. Two types of counters 2. MapReduce Join Chapter 8: MapReduce Hadoop Program To Join Data Chapter 9: Flume and Sqoop 1. What is SOOQP in Hadoop? 2. What is FLUME in Hadoop? 3. Some Important features of FLUME Chapter 10: Pig 1. Introduction to PIG 2. Create your First PIG Program 3. PART 1) Pig Installation 4. PART 2) Pig Demo Chapter 11: OOZIE 1. What is OOZIE? 2. How does OOZIE work? 3. Example Workflow Diagram 4. Oozie workflow application 5. Why use Oozie? 6. FEATURES OF OOZIE

Over 70 recipes to help you use Apache Spark as your single big data computing platform and master its libraries About This Book This book contains recipes on how to use Apache Spark as a unified compute engine Cover how to connect various source systems to Apache Spark Covers various parts of machine learning including supervised/unsupervised learning & recommendation engines Who This Book Is For This book is for data engineers, data scientists, and those who want to implement Spark for real-time data processing. Anyone who is using Spark (or is planning to) will benefit from this book. The book assumes you have a basic knowledge of Scala as a programming language. What You Will Learn Install and configure Apache Spark with various cluster managers & on AWS Set up a development environment for Apache Spark including Databricks Cloud notebook Find out how to operate on data in Spark with schemas Get to grips with real-time streaming analytics using Spark Streaming & Structured Streaming Master supervised learning and unsupervised learning using MLlib Build a recommendation engine using MLlib Graph processing using GraphX & GraphFrames libraries Develop a set of common applications or project types, and solutions that solve complex big data problems In Detail While Apache Spark 1.x gained a lot of traction and adoption in the early years, Spark 2.x delivers notable improvements in the areas of API, schema awareness, Performance, Structured Streaming, and simplifying building blocks to build better, faster, smarter, and more accessible big data applications. This book uncovers all these features in the form of structured recipes to analyze and mature large and complex sets of data. Starting with installing and configuring Apache Spark with various cluster managers, you will learn to set up development environments. Further on, you will be introduced to working with RDDs, DataFrames and Datasets to operate on schema aware data, and real-time streaming with various sources such as Twitter Stream and Apache Kafka. You will also work through recipes on machine learning, including supervised learning, unsupervised learning & recommendation engines in Spark. Last but not least, the final few chapters delve deeper into the concepts of graph processing using GraphX, securing your implementations, cluster optimization, and troubleshooting. Style and approach This book is packed with intuitive recipes supported with line-by-line explanations to help you understand Spark 2.x's real-time processing capabilities and deploy scalable big data solutions. This is a valuable resource for data scientists and those working on large-scale data projects.

Counsels programmers and administrators for big and small organizations on how to work with large-scale application datasets using Apache Hadoop, discussing its capacity for storing and processing large amounts of data while demonstrating best practices for building reliable and scalable distributed systems. 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

Data-intensive Text Processing with MapReduce

Agile Data Science 2.0

Practical Graph Analytics with Apache Giraph

Data Algorithms

Hadoop 2 Quick-Start Guide

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 Hadoop applications in production. 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.

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 a concise guide to the available MapReduce 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.

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 a concise guide to the available MapReduce 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.

Perform forensic investigations on Hadoop clusters with cutting-edge tools and techniques About This Book Identify, collect, and analyze Hadoop evidence forensically Learn about Hadoop's internals and Big Data file storage concepts A step-by-step guide to help you perform forensic analysis using freely available tools Who This Book Is For This book is meant for statisticians and forensic analysts with basic knowledge of digital forensics. They do not need to know Big Data Forensics. If you are an IT professional, law enforcement professional, legal professional, or a student interested in Big Data and forensics, this book is the perfect hands-on guide for learning how to conduct Hadoop forensic investigations. Each topic and step in the forensic process is described in accessible language. What You Will Learn Understand Hadoop internals and file storage Collect and analyze Hadoop forensic evidence Perform complex forensic analysis for fraud and other investigations Use state-of-the-art forensic tools Conduct interviews to identify Hadoop evidence Create compelling presentations of your forensic findings Understand how Big Data clusters operate Apply advanced forensic techniques in an investigation, including file carving, statistical analysis, and more In Detail Big Data forensics is an important type of digital investigation that involves the identification, collection, and analysis of large-scale Big Data systems. Hadoop is one of the most popular Big Data solutions, and forensically investigating a Hadoop cluster requires specialized tools and techniques. With the explosion of Big Data, forensic investigators need to be prepared to analyze the petabytes of data stored in Hadoop clusters. Understanding Hadoop's operational structure and performing forensic analysis with court-accepted tools and best practices will help you conduct a successful investigation. Discover how to perform a complete forensic investigation of large-scale Hadoop clusters using the same tools and techniques employed by forensic experts. This book begins by taking you through the process of forensic investigation and the pitfalls to avoid. It will walk you through Hadoop's internals and architecture, and you will discover what types of information Hadoop stores and how to access that data. You will learn to identify Big Data evidence using techniques to survey a live system and interview witnesses. After setting up your own Hadoop test environment, you will discover using techniques such as forensic imaging and application-based extraction. You will analyze Hadoop evidence using advanced tools and techniques to uncover events and statistical information. Finally, data visualization and evidence presentation techniques are covered to help you properly communicate your findings to any audience. Style and approach This book is a complete guide that follows every step of the forensic analysis process in detail. You will be guided through each key topic and step necessary to perform an investigation. Hands-on exercises are presented throughout the book, and technical reference guides and sample documents are included for real-world use.

This book is useful for Hadoop administrators who need to learn how to monitor and diagnose their clusters. Also, the book will prove useful for new users of the technology, as the language used is simple and easy to grasp.

Recipes for Scaling Up with Hadoop and Spark

Mastering Apache Cassandra - Second Edition

BARC Computer Science (CS/IT) Exam Prep Book 2022 | 10 Full-length Mock Tests (1000+ Solved Questions)

Learn the Essentials of Big Data Computing in the Apache Hadoop 2 Ecosystem

**Intelligent Systems Design and Applications**  
Intelligent Systems Design and Applications Excel 2013 to transform raw data into ready to use information About This Book Create formulas to help you analyze and explain findings Develop interactive spreadsheets that will impress your audience and give them the ability to slice and dice data A step-by-step guide to learn various ways to model data for businesses with the help of Excel Analyze and explain findings using the Analyze tool in Excel 2013 to transform raw data into ready to use information About This Book Create formulas to help you analyze and explain findings Develop interactive spreadsheets that will impress your audience and give them the ability to slice and dice data A step-by-step guide to learn various ways to model data for businesses with the help of Excel What you will start using in Excel 2013 for data analysis and business modeling and enhance your skills in the data analysis life cycle then this book is for you, whether you're new to Excel or experienced. What You Will Learn Discover what Excel formulas are all about and how to use them in your spreadsheet development Identify bad data and learn cleaning strategies Create interactive spreadsheets to leverage Excel's powerful built-in tools to get the median, maximum, and minimum values of your data Build impressive tables and combine datasets using Excel's built-in functionality Learn the powerful scripting language VBA, allowing you to implement your own custom solutions with ease In Detail Excel 2013 is one of the easiest to use data analysis tools you will ever come across. It has made it the go to tool for all your data needs. Complex operations with Excel, such as creating charts and graphs, visualization, and analyzing data make it a great tool for managers, data scientists, financial data analysts, and those who work closely with data. Learning data analysis and will help you bring your data skills to the next level. This book starts by walking you through creating your own spreadsheets and then the most widely used open source Hadoop ecosystem. You'll learn to discover how to spot bad data and strategies to clean that data to make it useful to you. Next, you'll learn to create custom columns, identify key metrics, and make decisions based on business rules. You'll create custom macros using new macros. Finally, at the end of the book, you'll be provided with useful shortcuts and tips, enabling you to do efficient data analysis and business modeling with Excel 2013. Style and approach This is a step-by-step guide to performing data analysis and business modeling with Excel 2013, complete with examples and tips.

Advanced Analytics with Apache Spark, Mahout and Hadoop 2. This book covers the advanced topics of Hadoop ecosystem, 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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 Hadoop cluster to efficiently process 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

*Big Data Forensics – Learning Hadoop Investigations*

*Hadoop Operations*

*Elasticsearch for Hadoop*

*Monitoring Hadoop*

*Learns Hadoop in 24 Hours*

The book is aimed at intermediate developers with an understanding of core database concepts who want to become a master at implementing Cassandra for their application.

MapReduce Design Patterns

Building Effective Algorithms and Analytics for Hadoop and Other Systems

Apache Hadoop YARN

Hadoop: The Definitive Guide

Apache Hadoop 3 Quick Start Guide