

20 Essential Hadoop Tools For Crunching Big Data

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.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book.

The Comprehensive, Up-to-Date Apache Hadoop Administration Handbook and Reference

“Sam Alapati has worked with production Hadoop clusters for six years. His unique depth of experience has enabled him to write the go-to resource for all administrators looking to spec, size, expand, and secure production Hadoop clusters of any size.” —Paul Dix, Series Editor

In Expert Hadoop® Administration, leading Hadoop administrator Sam R. Alapati brings together authoritative knowledge for creating, configuring, securing, managing, and optimizing production Hadoop clusters in any environment. Drawing on his experience with large-scale Hadoop administration, Alapati integrates action-oriented advice with carefully researched explanations of both problems and solutions. He covers an unmatched range of topics and offers an unparalleled collection of realistic examples. Alapati demystifies complex Hadoop environments, helping you understand exactly what happens behind the scenes when you administer your cluster. You’ll gain unprecedented insight as you walk through building clusters from scratch and configuring high availability, performance, security, encryption, and other key attributes. The high-value administration skills you learn here will be indispensable no matter what Hadoop distribution you use or what Hadoop applications you run.

Understand Hadoop’s architecture from an administrator’s standpoint

- Create simple and fully distributed clusters
- Run MapReduce and Spark applications in a Hadoop cluster
- Manage and protect Hadoop data and high availability
- Work with HDFS commands, file permissions, and storage management
- Move data, and use

YARN to allocate resources and schedule jobs Manage job workflows with Oozie and Hue Secure, monitor, log, and optimize Hadoop Benchmark and troubleshoot Hadoop

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

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

Spanning the multi-disciplinary scope of information technology, the Encyclopedia of Information Systems and Technology draws together comprehensive coverage of the inter-related aspects of information systems and technology. The topics covered in this encyclopedia encompass internationally recognized bodies of knowledge, including those of The IT BOK, the Chartered Information Technology Professionals Program, the International IT Professional Practice Program (British Computer Society), the Core Body of Knowledge for IT Professionals (Australian Computer Society), the International Computer Driving License Foundation (European Computer Driving License Foundation), and the Guide to the Software Engineering Body of Knowledge. Using the universally recognized definitions of IT and information systems from these recognized bodies of knowledge, the encyclopedia brings together the information that students, practicing professionals, researchers, and academicians need to keep their knowledge up to date. Also Available Online This Taylor & Francis encyclopedia is also available through online subscription, offering a variety of extra benefits for researchers, students, and librarians, including: ? Citation tracking and alerts ? Active reference linking ? Saved searches and marked lists ? HTML and PDF format options Contact Taylor and Francis for more information or to inquire about subscription options and print/online combination packages. US: (Tel) 1.888.318.2367; (E-mail) e-reference@taylorandfrancis.com International: (Tel) +44 (0) 20 7017 6062; (E-mail) online.sales@tandf.co.uk

The idea of the 1st International Conference on Intelligent Computing and Applications (ICICA 2014) is to bring the Research Engineers, Scientists, Industrialists, Scholars and Students together from in and around the globe to present the on-going research activities and hence to encourage research interactions between universities and industries. The conference provides opportunities for the delegates to exchange new ideas, applications and experiences, to establish research relations and to find global partners for future collaboration. The proceedings covers latest progresses in the cutting-edge research on various research areas of Image, Language Processing, Computer Vision and Pattern Recognition, Machine Learning, Data Mining and Computational Life Sciences, Management of Data including Big Data and Analytics, Distributed and Mobile Systems including Grid and Cloud infrastructure, Information Security and Privacy, VLSI, Electronic Circuits, Power Systems, Antenna, Computational fluid dynamics & Heat transfer, Intelligent Manufacturing, Signal Processing, Intelligent Computing, Soft Computing, Bio-informatics, Bio Computing, Web Security, Privacy and E-Commerce, E-governance, Service Orient Architecture, Data Engineering, Open Systems, Optimization, Communications, Smart wireless and sensor Networks, Smart Antennae, Networking and Information security, Machine Learning, Mobile Computing and Applications, Industrial Automation and MES, Cloud Computing, Green IT, IT for Rural Engineering, Business Computing, Business Intelligence, ICT for Education for solving hard problems, and finally to create awareness about these domains to a wider audience of practitioners.

This book constitutes the proceedings of the 9th International Conference on Big Data, BigData 2020, held as part of SCF 2020, during September 18-20, 2020. The conference was planned to take place in Honolulu, HI, USA and was changed to a virtual format due to the COVID-19 pandemic. The 16 full and 3 short papers presented were carefully reviewed and selected from 52 submissions. The topics covered are Big Data Architecture, Big Data Modeling, Big Data As A Service, Big Data for Vertical Industries (Government, Healthcare, etc.), Big Data Analytics, Big Data Toolkits, Big Data Open Platforms, Economic Analysis, Big Data for Enterprise Transformation, Big Data in Business Performance Management, Big Data for Business Model Innovations and Analytics, Big Data in Enterprise Management Models and Practices, Big Data in Government Management Models and Practices, and Big Data in Smart Planet Solutions.

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. This book is also meant for Hadoop professionals who want to find solutions to the different challenges they come across in their Hadoop projects.

Learn how to use the Apache Hadoop projects, including MapReduce, HDFS, Apache Hive, Apache HBase, Apache Kafka, Apache Mahout, and Apache Solr. From setting up the environment to running sample applications each chapter in

this book is a practical tutorial on using an Apache Hadoop ecosystem project. While several books on Apache Hadoop are available, most are based on the main projects, MapReduce and HDFS, and none discusses the other Apache Hadoop ecosystem projects and how they all work together as a cohesive big data development platform. What You Will Learn: Set up the environment in Linux for Hadoop projects using Cloudera Hadoop Distribution CDH 5 Run a MapReduce job Store data with Apache Hive, and Apache HBase Index data in HDFS with Apache Solr Develop a Kafka messaging system Stream Logs to HDFS with Apache Flume Transfer data from MySQL database to Hive, HDFS, and HBase with Sqoop Create a Hive table over Apache Solr Develop a Mahout User Recommender System Who This Book Is For: Apache Hadoop developers. Pre-requisite knowledge of Linux and some knowledge of Hadoop is required. Although there are already some books published on Big Data, most of them only cover basic concepts and society impacts and ignore the internal implementation details-making them unsuitable to R&D people. To fill such a need, Big Data: Storage, Sharing, and Security examines Big Data management from an R&D perspective. It covers the 3S desi

This book presents a compilation of current trends, technologies, and challenges in connection with Big Data. Many fields of science and engineering are data-driven, or generate huge amounts of data that are ripe for the picking. There are now more sources of data than ever before, and more means of capturing data. At the same time, the sheer volume and complexity of the data have sparked new developments, where many Big Data problems require new solutions. Given its scope, the book offers a valuable reference guide for all graduate students, researchers, and scientists interested in exploring the potential of Big Data applications.

Harness the power of Python to develop data mining applications, analyze data, delve into machine learning, explore object detection using Deep Neural Networks, and create insightful predictive models. About This Book Use a wide variety of Python libraries for practical data mining purposes. Learn how to find, manipulate, analyze, and visualize data using Python. Step-by-step instructions on data mining techniques with Python that have real-world applications. Who This Book Is For If you are a Python programmer who wants to get started with data mining, then this book is for you. If you are a data analyst who wants to leverage the power of Python to perform data mining efficiently, this book will also help you. No previous experience with data mining is expected. What You Will Learn Apply data mining concepts to real-world problems Predict the outcome of sports matches based on past results Determine the author of a document based on their writing style Use APIs to download datasets from social media and other online services Find and extract good features from difficult datasets Create models that solve real-world problems Design and develop data mining applications using a variety of datasets Perform object detection in images using Deep Neural Networks Find meaningful insights from your data

through intuitive visualizations Compute on big data, including real-time data from the internet In Detail This book teaches you to design and develop data mining applications using a variety of datasets, starting with basic classification and affinity analysis. This book covers a large number of libraries available in Python, including the Jupyter Notebook, pandas, scikit-learn, and NLTK. You will gain hands on experience with complex data types including text, images, and graphs. You will also discover object detection using Deep Neural Networks, which is one of the big, difficult areas of machine learning right now. With restructured examples and code samples updated for the latest edition of Python, each chapter of this book introduces you to new algorithms and techniques. By the end of the book, you will have great insights into using Python for data mining and understanding of the algorithms as well as implementations. Style and approach This book will be your comprehensive guide to learning the various data mining techniques and implementing them in Python. A variety of real-world datasets is used to explain data mining techniques in a very crisp and easy to understand manner.

This book takes you on a fantastic journey to discover the attributes of big data using Apache Hive. Key Features Grasp the skills needed to write efficient Hive queries to analyze the Big Data Discover how Hive can coexist and work with other tools within the Hadoop ecosystem Uses practical, example-oriented scenarios to cover all the newly released features of Apache Hive 2.3.3 Book Description In this book, we prepare you for your journey into big data by firstly introducing you to backgrounds in the big data domain, alongwith the process of setting up and getting familiar with your Hive working environment. Next, the book guides you through discovering and transforming the values of big data with the help of examples. It also hones your skills in using the Hive language in an efficient manner. Toward the end, the book focuses on advanced topics, such as performance, security, and extensions in Hive, which will guide you on exciting adventures on this worthwhile big data journey. By the end of the book, you will be familiar with Hive and able to work effeciently to find solutions to big data problems What you will learn Create and set up the Hive environment Discover how to use Hive's definition language to describe data Discover interesting data by joining and filtering datasets in Hive Transform data by using Hive sorting, ordering, and functions Aggregate and sample data in different ways Boost Hive query performance and enhance data security in Hive Customize Hive to your needs by using user-defined functions and integrate it with other tools Who this book is for If you are a data analyst, developer, or simply someone who wants to quickly get started with Hive to explore and analyze Big Data in Hadoop, this is the book for you. Since Hive is an SQL-like language, some previous experience with SQL will be useful to get the most out of this book.

This two-volume set LNCS 6691 and 6692 constitutes the refereed proceedings of the 11th International Work-Conference on Artificial Neural Networks, IWANN 2011, held in Torremolinos-Málaga, Spain, in June 2011. The 154 revised papers

were carefully reviewed and selected from 202 submissions for presentation in two volumes. The first volume includes 69 papers organized in topical sections on mathematical and theoretical methods in computational intelligence; learning and adaptation; bio-inspired systems and neuro-engineering; hybrid intelligent systems; applications of computational intelligence; new applications of brain-computer interfaces; optimization algorithms in graphic processing units; computing languages with bio-inspired devices and multi-agent systems; computational intelligence in multimedia processing; and biologically plausible spiking neural processing.

This in-depth guide provides managers with a solid understanding of data and data trends, the opportunities that it can offer to businesses, and the dangers of these technologies. Written in an accessible style, Steven Finlay provides a contextual roadmap for developing solutions that deliver benefits to organizations.

The digital age has presented an exponential growth in the amount of data available to individuals looking to draw conclusions based on given or collected information across industries. Challenges associated with the analysis, security, sharing, storage, and visualization of large and complex data sets continue to plague data scientists and analysts alike as traditional data processing applications struggle to adequately manage big data. The Handbook of Research on Big Data Storage and Visualization Techniques is a critical scholarly resource that explores big data analytics and technologies and their role in developing a broad understanding of issues pertaining to the use of big data in multidisciplinary fields. Featuring coverage on a broad range of topics, such as architecture patterns, programming systems, and computational energy, this publication is geared towards professionals, researchers, and students seeking current research and application topics on the subject.

This IBM® Redbooks® publication documents how IBM Platform Computing, with its IBM Platform Symphony® MapReduce framework, IBM Spectrum Scale (based Upon IBM GPFS™), IBM Platform LSF®, the Advanced Service Controller for Platform Symphony are work together as an infrastructure to manage not just Hadoop-related offerings, but many popular industry offerings such as Apache Spark, Storm, MongoDB, Cassandra, and so on. It describes the different ways to run Hadoop in a big data environment, and demonstrates how IBM Platform Computing solutions, such as Platform Symphony and Platform LSF with its MapReduce Accelerator, can help performance and agility to run Hadoop on distributed workload managers offered by IBM. This information is for technical professionals (consultants, technical support staff, IT architects, and IT specialists) who are responsible for delivering cost-effective cloud services and big data solutions on IBM Power Systems™ to help uncover insights among client's data so they can optimize product development and business results. 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.

Conference Proceedings of 20th European Conference on Research Methods in Business and Management

You've heard the hype about Hadoop: it runs petabyte-scale data mining tasks insanely fast, it runs gigantic tasks on clouds for absurdly cheap, it's been heavily committed to by tech giants like IBM, Yahoo!, and the Apache Project, and it's completely open-source (thus free). But what exactly is it, and more importantly, how do you even get a Hadoop cluster up and running? From Apress, the name you've come to trust for hands-on technical knowledge, Pro Hadoop brings you up to speed on Hadoop. You learn the ins and outs of MapReduce; how to structure a cluster, design, and implement the Hadoop file system; and how to

build your first cloud-computing tasks using Hadoop. Learn how to let Hadoop take care of distributing and parallelizing your software—you just focus on the code, Hadoop takes care of the rest. Best of all, you'll learn from a tech professional who's been in the Hadoop scene since day one. Written from the perspective of a principal engineer with down-in-the-trenches knowledge of what to do wrong with Hadoop, you learn how to avoid the common, expensive first errors that everyone makes with creating their own Hadoop system or inheriting someone else's. Skip the novice stage and the expensive, hard-to-fix mistakes...go straight to seasoned pro on the hottest cloud-computing framework with Pro Hadoop. Your productivity will blow your managers away.

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures Get expert guidance on architecting end-to-end data management solutions with Apache Hadoop. While many sources explain how to use various components in the Hadoop ecosystem, this practical book takes you through architectural considerations necessary to tie those components together into a complete tailored application, based on your particular use case. To reinforce those lessons, the book's second section provides detailed examples of architectures used in some of the most commonly found Hadoop applications. Whether you're designing a new Hadoop application, or planning to integrate Hadoop into your existing data infrastructure, Hadoop Application Architectures will skillfully guide you through the process. This book covers: Factors to consider when using Hadoop to store and model data Best practices for moving data in and out of the system Data processing frameworks, including MapReduce, Spark, and Hive Common Hadoop processing patterns, such as removing duplicate records and using windowing analytics Giraph, GraphX, and other tools for large graph processing on Hadoop Using workflow orchestration and scheduling tools such as Apache Oozie Near-real-time stream processing with Apache Storm, Apache

Spark Streaming, and Apache Flume Architecture examples for clickstream analysis, fraud detection, and data warehousing

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 system, you will collect evidence using techniques such as forensic imaging and application-based extractions. 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.

Health Informatics: Practical Guide focuses on the application of information technology in healthcare to improve individual and population health, education and research. The goal of the seventh edition is to stimulate and educate healthcare and IT professionals and students about the key topics in this rapidly changing field. Dr. William Hersh from Oregon Health & Science University is the co-editor and author of multiple chapters. Topics include Health Informatics (HI) overview, electronic health records, healthcare data analytics, health information exchange, architecture of information systems, evidence-based medicine, consumer health informatics, HI ethics, quality improvement strategies and more. The 22 chapters feature learning objectives, case studies, recommended reading, future trends, key points, conclusions and over 1800 references. It is available as a paperback and an eBook. Visit the textbook companion website at <http://informaticseducation.org/> for more information.

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

Get Started Fast with Apache Hadoop® 2, YARN, and Today's Hadoop Ecosystem With Hadoop 2.x and YARN, Hadoop moves beyond MapReduce to become practical for virtually any type of data processing. Hadoop 2.x and the Data Lake concept represent a radical shift away from conventional approaches to data usage and storage. Hadoop 2.x installations offer unmatched scalability and breakthrough extensibility that supports new and existing Big Data analytics processing methods and models. **Hadoop® 2 Quick-Start Guide** is the first easy, accessible guide to Apache Hadoop 2.x, YARN, and the modern Hadoop ecosystem. Building on his unsurpassed experience teaching Hadoop and Big Data, author Douglas Eadline covers all the basics you need to know to install and use Hadoop 2 on personal computers or servers, and to navigate the powerful technologies that complement it. Eadline concisely introduces and explains every key Hadoop 2 concept, tool, and service, illustrating each with a simple “beginning-to-end” example and identifying trustworthy, up-to-date resources for learning more. This guide is ideal if you want to learn about Hadoop

2 without getting mired in technical details. Douglas Eadline will bring you up to speed quickly, whether you're a user, admin, devops specialist, programmer, architect, analyst, or data scientist. Coverage Includes Understanding what Hadoop 2 and YARN do, and how they improve on Hadoop 1 with MapReduce Understanding Hadoop-based Data Lakes versus RDBMS Data Warehouses Installing Hadoop 2 and core services on Linux machines, virtualized sandboxes, or clusters Exploring the Hadoop Distributed File System (HDFS) Understanding the essentials of MapReduce and YARN application programming Simplifying programming and data movement with Apache Pig, Hive, Sqoop, Flume, Oozie, and HBase Observing application progress, controlling jobs, and managing workflows Managing Hadoop efficiently with Apache Ambari—including recipes for HDFS to NFSv3 gateway, HDFS snapshots, and YARN configuration Learning basic Hadoop 2 troubleshooting, and installing Apache Hue and Apache Spark Offers information on the architecture and data design necessary to create Hadoop-based enterprise applications.

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 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.

Integrate, deploy, rapidly configure, and successfully manage your own big data-intensive clusters in the cloud using OpenStack Sahara About This Book A fast paced guide to help you utilize the benefits of Sahara in OpenStack to meet the Big Data world of Hadoop. A step by step approach to simplify the complexity of Hadoop configuration, deployment and maintenance. Who This Book Is For This book targets data scientists, cloud developers and Devops Engineers who would like to become proficient with OpenStack Sahara. Ideally, this book is well suitable for readers who are familiar with databases, Hadoop and Spark solutions. Additionally, a basic prior knowledge of OpenStack is expected. The readers should also be familiar with different Linux boxes, distributions and virtualization technology. What You Will Learn Integrate and Install Sahara with OpenStack environment Learn Sahara architecture under the hood Rapidly configure and scale Hadoop clusters on top of OpenStack Explore the Sahara REST API to create, deploy and manage a Hadoop cluster Learn the Elastic Processing Data (EDP) facility to execute jobs in clusters from Sahara Cover other Hadoop stable plugins existing supported by Sahara Discover different features provided by Sahara for Hadoop provisioning and deployment Learn how to troubleshoot OpenStack Sahara issues In Detail The Sahara project is a module that aims to simplify the building of data processing capabilities on OpenStack. The goal of this book is to provide a focused, fast paced guide to installing, configuring, and getting started with integrating Hadoop with OpenStack, using Sahara. The book should explain to users how to deploy their data-intensive Hadoop and Spark clusters on top of OpenStack. It will also cover how to use the Sahara REST API, how to develop applications for Elastic Data Processing on Openstack, and setting up hadoop or spark clusters on Openstack. Style and approach This book takes a step by step approach teaching how to integrate, deploy and manage data using OpenStack Sahara. It will teach how the OpenStack Sahara is beneficial by simplifying the complexity of Hadoop configuration, deployment and maintenance.

We have never lived at a time of faster and more transformative technological and societal changes. It can be hard for executives to keep up with the developments and shifts. This book cuts through all of the hype and presents the key business trends anyone should be aware of now as they will shape businesses into the foreseeable future. Business Trends in Practice includes case studies across all industries, with companies such as: Tesla, Ocado, Netflix, Microsoft, Google, Alibaba, Rolls Royce, Mercedes Benz, Apple, and many more. Some of the key trends the author will examine include: The AI revolution Robots

and business processes automation Remote working, working from home and new flexibility Social & environmental Responsibility Increased Diversity As part of Bernard Marr's popular 'In Practice' series, Business Trends in Practice will help you identify the key business trends that will keep you one step ahead of the competition.

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.

Residents in Boston, Massachusetts are automatically reporting potholes and road hazards via their smartphones. Progressive Insurance tracks real-time customer driving patterns and uses that information to offer rates truly commensurate with individual safety. Google accurately predicts local flu outbreaks based upon thousands of user search queries. Amazon provides remarkably insightful, relevant, and timely product recommendations to its hundreds of millions of customers. Quantcast lets companies target precise audiences and key demographics throughout the Web. NASA runs contests via gamification site TopCoder, awarding prizes to those with the most innovative and cost-effective solutions to its problems. Explorys offers penetrating and previously unknown insights into healthcare behavior. How do these organizations and municipalities do it? Technology is certainly a big part, but in each case the answer lies deeper than that. Individuals at these organizations have realized that they don't have to be Nate Silver to reap massive benefits from today's new and emerging types of data. And each of these organizations has embraced Big Data, allowing them to make astute and otherwise impossible observations, actions, and predictions. It's time to start thinking big. In *Too Big to Ignore*, recognized technology expert and award-winning author Phil Simon explores an unassailably important trend: Big Data, the massive amounts, new types, and multifaceted sources of information streaming at us faster than ever. Never before have we seen data with the volume, velocity, and variety of today. Big Data is no temporary blip of fad. In fact, it is only going to intensify in the coming years, and its ramifications for the future of business are impossible to overstate. *Too Big to Ignore* explains why Big Data is a big deal. Simon provides commonsense, jargon-free advice for people and organizations looking to understand and leverage Big Data. Rife with case studies, examples, analysis, and quotes from real-world Big Data practitioners, the book is required reading for chief executives, company owners, industry leaders, and business professionals.

Hadoop Essentials Packt Publishing Ltd

Society is now completely driven by data with many industries relying on data to conduct business or basic functions within the organization. With the efficiencies that big data bring to all institutions, data is continuously being collected and analyzed. However, data sets may be

too complex for traditional data-processing, and therefore, different strategies must evolve to solve the issue. The field of big data works as a valuable tool for many different industries. The Research Anthology on Big Data Analytics, Architectures, and Applications is a complete reference source on big data analytics that offers the latest, innovative architectures and frameworks and explores a variety of applications within various industries. Offering an international perspective, the applications discussed within this anthology feature global representation. Covering topics such as advertising curricula, driven supply chain, and smart cities, this research anthology is ideal for data scientists, data analysts, computer engineers, software engineers, technologists, government officials, managers, CEOs, professors, graduate students, researchers, and academicians.

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

This comprehensive edited volume is the first of its kind, designed to serve as a textbook for long-duration business analytics programs. It can also be used as a guide to the field by practitioners. The book has contributions from experts in top universities and industry. The editors have taken extreme care to ensure continuity across the chapters. The material is organized into three parts: A) Tools, B) Models and C) Applications. In Part A, the tools used by business analysts are described in detail. In Part B, these tools are applied to construct models used to solve business problems. Part C contains detailed applications in various functional areas of business and several case studies. Supporting material can be found in the appendices that develop the pre-requisites for the main text. Every chapter has a business orientation. Typically, each chapter begins with the description of business problems that are transformed into data questions; and methodology is developed to solve these questions. Data analysis is conducted using widely used software, the output and results are clearly explained at each stage of development. These are finally transformed into a business solution. The

companion website provides examples, data sets and sample code for each chapter.

Do you want to broaden your Hadoop skill set and take your knowledge to the next level? Do you wish to enhance your knowledge of Hadoop to solve challenging data processing problems? Are your Hadoop jobs, Pig scripts, or Hive queries not working as fast as you intend? Are you looking to understand the benefits of upgrading Hadoop? If the answer is yes to any of these, this book is for you. It assumes novice-level familiarity with Hadoop.

This book constitutes the proceedings of the Second International Conference on Information and Communication Technology for Development for Africa, ICT4DA 2019, held in Bahir Dar, Ethiopia, in May 2019. The 29 revised full papers presented were carefully reviewed and selected from 69 submissions. The papers address the impact of ICT in fostering economic development in Africa. In detail they cover the following topics: artificial intelligence and data science; wireless and mobile computing; and Natural Language Processing.

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.

This book includes selected papers from the International Conference on Data Science and Intelligent Applications (ICDSIA 2020), hosted by Gandhinagar Institute of Technology (GIT), Gujarat, India, on January 24–25, 2020. The proceedings present original and high-quality contributions on theory and practice concerning emerging technologies in the areas of data science and intelligent applications. The conference provides a forum for researchers from academia and industry to present and share their ideas, views and results, while also helping them approach the challenges of technological advancements from different viewpoints. The contributions cover a broad range of topics, including: collective intelligence, intelligent systems, IoT, fuzzy systems, Bayesian networks, ant colony optimization, data privacy and security, data mining, data warehousing, big data analytics, cloud computing, natural language processing, swarm intelligence, speech processing, machine learning and deep learning, and intelligent applications and systems. Helping strengthen the links between academia and industry, the book offers a valuable resource for instructors, students, industry practitioners, engineers, managers, researchers, and scientists alike.

[Copyright: 794dfa4f27c27640542154e5a1c73a3a](https://www.manning.com/books/hadoop-in-action)