

Practical Hadoop By Example

Explore big data concepts, platforms, analytics, and their applications using the power of Hadoop 3 Key Features Learn Hadoop 3 to build effective big data analytics solutions on-premise and on cloud Integrate Hadoop with other big data tools such as R, Python, Apache Spark, and Apache Flink Exploit big data using Hadoop 3 with real-world examples Book Description Apache Hadoop is the most popular platform for big data processing, and can be combined with a host of other big data tools to build powerful analytics solutions. Big Data Analytics with Hadoop 3 shows you how to do just that, by providing insights into the software as well as its benefits with the help of practical examples. Once you have taken a tour of Hadoop 3's latest features, you will get an overview of HDFS, MapReduce, and YARN, and how they enable faster, more efficient big data processing. You will then move on to learning how to integrate Hadoop with the open source tools, such as Python and R, to analyze and visualize data and perform statistical computing on big data. As you get acquainted with all this, you will explore how to use Hadoop 3 with Apache Spark and Apache Flink for real-time data analytics and stream processing. In addition to this, you will understand how to use Hadoop to build analytics solutions on the cloud and an end-to-end pipeline to perform big data analysis using practical use cases. By the end of this book, you will be well-versed with the analytical capabilities of the Hadoop ecosystem. You will be able to build powerful solutions to perform big data analytics and get insight effortlessly.

Get Free Practical Hadoop By Example

What you will learn Explore the new features of Hadoop 3 along with HDFS, YARN, and MapReduce Get well-versed with the analytical capabilities of Hadoop ecosystem using practical examples Integrate Hadoop with R and Python for more efficient big data processing Learn to use Hadoop with Apache Spark and Apache Flink for real-time data analytics Set up a Hadoop cluster on AWS cloud Perform big data analytics on AWS using Elastic Map Reduce Who this book is for Big Data Analytics with Hadoop 3 is for you if you are looking to build high-performance analytics solutions for your enterprise or business using Hadoop 3's powerful features, or you're new to big data analytics. A basic understanding of the Java programming language is required.

Ready to use statistical and machine-learning techniques across large data sets? This practical guide shows you why the Hadoop ecosystem is perfect for the job. Instead of deployment, operations, or software development usually associated with distributed computing, you'll focus on particular analyses you can build, the data warehousing techniques that Hadoop provides, and higher order data workflows this framework can produce. Data scientists and analysts will learn how to perform a wide range of techniques, from writing MapReduce and Spark applications with Python to using advanced modeling and data management with Spark MLlib, Hive, and HBase. You'll also learn about the analytical processes and data systems available to build and empower data products that can handle—and actually require—huge amounts of data. Understand core

Get Free Practical Hadoop By Example

concepts behind Hadoop and cluster computing Use design patterns and parallel analytical algorithms to create distributed data analysis jobs Learn about data management, mining, and warehousing in a distributed context using Apache Hive and HBase Use Sqoop and Apache Flume to ingest data from relational databases Program complex Hadoop and Spark applications with Apache Pig and Spark DataFrames Perform machine learning techniques such as classification, clustering, and collaborative filtering with Spark's MLlib

The go-to guidebook for deploying Big Data solutions with Hadoop Today's enterprise architects need to understand how the Hadoop frameworks and APIs fit together, and how they can be integrated to deliver real-world solutions. This book is a practical, detailed guide to building and implementing those solutions, with code-level instruction in the popular Wrox tradition. It covers storing data with HDFS and Hbase, processing data with MapReduce, and automating data processing with Oozie. Hadoop security, running Hadoop with Amazon Web Services, best practices, and automating Hadoop processes in real time are also covered in depth. With in-depth code examples in Java and XML and the latest on recent additions to the Hadoop ecosystem, this complete resource also covers the use of APIs, exposing their inner workings and allowing architects and developers to better leverage and customize them. The ultimate guide for developers, designers, and architects who need to build and deploy Hadoop applications Covers storing and processing data with various technologies, automating data processing,

Get Free Practical Hadoop By Example

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.

The book contains the latest trend in IT industry 'Big Data and Hadoop'. It explains how big is 'Big Data' and why everybody is trying to implement this into their IT project. It includes research work on various topics, theoretical and practical approach, each component of the architecture is described along with current industry trends. Big Data and Hadoop have taken together are a new skill as per the industry standards. Readers will get a compact book along with the industry experience and would be a reference to help readers.

KEY FEATURES
Overview Of Big Data, Basics of Hadoop, Hadoop Distributed File System, HBase, MapReduce, HIVE: The Data Warehouse Of Hadoop, PIG: The Higher Level Programming Environment, SQOOP: Importing Data From Heterogeneous Sources, Flume, Ozzie, Zookeeper & Big Data Stream Mining, Chapter-wise Questions & Previous Years Questions

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

Get Free Practical Hadoop By Example

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

Get Free Practical Hadoop By Example

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

Tackle the real-world complexities of modern machine learning with innovative, cutting-edge, techniques About This Book Fully-coded working examples using a wide range of machine learning libraries and tools, including Python, R, Julia, and Spark Comprehensive practical solutions taking you into the future of machine learning Go a step further and integrate your machine learning projects with Hadoop Who This Book Is For This book has been created for data scientists who want to see machine learning in action and explore its real-world application. With guidance on everything from the fundamentals of machine learning and predictive analytics to the latest innovations set to lead the big data revolution into the future, this is an unmissable resource for anyone dedicated to tackling current big data challenges. Knowledge of programming (Python and R) and mathematics is advisable if you want to get started immediately. What You Will Learn Implement a wide range of algorithms and techniques for tackling complex data Get to grips with some of the most powerful languages in data science, including R, Python, and Julia Harness the capabilities of Spark and Hadoop to manage and process data successfully Apply the

Get Free Practical Hadoop By Example

appropriate machine learning technique to address real-world problems Get acquainted with Deep learning and find out how neural networks are being used at the cutting-edge of machine learning Explore the future of machine learning and dive deeper into polyglot persistence, semantic data, and more In Detail Finding meaning in increasingly larger and more complex datasets is a growing demand of the modern world. Machine learning and predictive analytics have become the most important approaches to uncover data gold mines. Machine learning uses complex algorithms to make improved predictions of outcomes based on historical patterns and the behaviour of data sets. Machine learning can deliver dynamic insights into trends, patterns, and relationships within data, immensely valuable to business growth and development. This book explores an extensive range of machine learning techniques uncovering hidden tricks and tips for several types of data using practical and real-world examples. While machine learning can be highly theoretical, this book offers a refreshing hands-on approach without losing sight of the underlying principles. Inside, a full exploration of the various algorithms gives you high-quality guidance so you can begin to see just how effective machine learning is at tackling contemporary challenges of big data. This is the only book you need to implement a whole suite of open source tools, frameworks, and languages in machine learning. We will cover the leading data science languages, Python and R, and the underrated but powerful Julia, as well as a range of other big data

Get Free Practical Hadoop By Example

platforms including Spark, Hadoop, and Mahout. Practical Machine Learning is an essential resource for the modern data scientists who want to get to grips with its real-world application. With this book, you will not only learn the fundamentals of machine learning but dive deep into the complexities of real world data before moving on to using Hadoop and its wider ecosystem of tools to process and manage your structured and unstructured data. You will explore different machine learning techniques for both supervised and unsupervised learning; from decision trees to Naive Bayes classifiers and linear and clustering methods, you will learn strategies for a truly advanced approach to the statistical analysis of data. The book also explores the cutting-edge advancements in machine learning, with worked examples and guidance on deep learning and reinforcement learning, providing you with practical demonstrations and samples that help take the theory—and mystery—out of even the most advanced machine learning methodologies. Style and approach A practical data science tutorial designed to give you an insight into the practical application of machine learning, this book takes you through complex concepts and tasks in an accessible way. Featuring information on a wide range of data science techniques, Practical Machine Learning is a comprehensive data science resource. Until recently, Hadoop deployments existed on hardware owned and run by organizations. Now, of course, you can acquire the computing resources and network connectivity to run Hadoop clusters in the cloud. But there's a lot more to deploying Hadoop to the public

Get Free Practical Hadoop By Example

cloud than simply renting machines. This hands-on guide shows developers and systems administrators familiar with Hadoop how to install, use, and manage cloud-born clusters efficiently. You'll learn how to architect clusters that work with cloud-provider features—not just to avoid pitfalls, but also to take full advantage of these services. You'll also compare the Amazon, Google, and Microsoft clouds, and learn how to set up clusters in each of them. Learn how Hadoop clusters run in the cloud, the problems they can help you solve, and their potential drawbacks Examine the common concepts of cloud providers, including compute capabilities, networking and security, and storage Build a functional Hadoop cluster on cloud infrastructure, and learn what the major providers require Explore use cases for high availability, relational data with Hive, and complex analytics with Spark Get patterns and practices for running cloud clusters, from designing for price and security to dealing with maintenance

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

Get Free Practical Hadoop By Example

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

The practical implications of technical debt for the entire software lifecycle; with examples and case studies. Technical debt in software is incurred when developers take shortcuts and make ill-advised technical decisions in the initial phases of a project, only to be confronted with the need for costly and labor-intensive workarounds later. This book offers advice on how to avoid technical debt, how to locate its sources, and how to remove it. It focuses on the practical implications of technical debt for the entire software life cycle, with examples and case studies from companies that range from Boeing to Twitter. Technical debt is normal; it is part of most iterative development processes. But if debt is ignored, over time it may become unmanageably complex, requiring developers to spend all of their effort fixing bugs, with no time to add new features--and after all,

Get Free Practical Hadoop By Example

new features are what customers really value. The authors explain how to monitor technical debt, how to measure it, and how and when to pay it down. Broadening the conventional definition of technical debt, they cover requirements debt, implementation debt, testing debt, architecture debt, documentation debt, deployment debt, and social debt. They intersperse technical discussions with "Voice of the Practitioner" sidebars that detail real-world experiences with a variety of technical debt issues.

Practical Business Analytics Using SAS: A Hands-on Guide shows SAS users and businesspeople how to analyze data effectively in real-life business scenarios. The book begins with an introduction to analytics, analytical tools, and SAS programming. The authors—both SAS, statistics, analytics, and big data experts—first show how SAS is used in business, and then how to get started programming in SAS by importing data and learning how to manipulate it. Besides illustrating SAS basic functions, you will see how each function can be used to get the information you need to improve business performance. Each chapter offers hands-on exercises drawn from real business situations. The book then provides an overview of statistics, as well as instruction on exploring data, preparing it for analysis, and testing hypotheses. You will learn how to use SAS to perform analytics and model using both basic and advanced techniques like multiple

Get Free Practical Hadoop By Example

regression, logistic regression, and time series analysis, among other topics. The book concludes with a chapter on analyzing big data. Illustrations from banking and other industries make the principles and methods come to life. Readers will find just enough theory to understand the practical examples and case studies, which cover all industries. Written for a corporate IT and programming audience that wants to upgrade skills or enter the analytics field, this book includes: More than 200 examples and exercises, including code and datasets for practice. Relevant examples for all industries. Case studies that show how to use SAS analytics to identify opportunities, solve complicated problems, and chart a course. Practical Business Analytics Using SAS: A Hands-on Guide gives you the tools you need to gain insight into the data at your fingertips, predict business conditions for better planning, and make excellent decisions. Whether you are in retail, finance, healthcare, manufacturing, government, or any other industry, this book will help your organization increase revenue, drive down costs, improve marketing, and satisfy customers better than ever before.

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

Get Free Practical Hadoop By Example

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.

Finding patterns in massive event streams can be difficult, but learning how to find them doesn't have to be. This unique hands-on guide shows you how to solve this and many other problems in large-scale data processing with simple, fun, and elegant tools that leverage Apache Hadoop. You'll gain a practical, actionable view of big data by working with real data and real problems. Perfect for beginners,

Get Free Practical Hadoop By Example

this book's approach will also appeal to experienced practitioners who want to brush up on their skills. Part I explains how Hadoop and MapReduce work, while Part II covers many analytic patterns you can use to process any data. As you work through several exercises, you'll also learn how to use Apache Pig to process data. Learn the necessary mechanics of working with Hadoop, including how data and computation move around the cluster Dive into map/reduce mechanics and build your first map/reduce job in Python Understand how to run chains of map/reduce jobs in the form of Pig scripts Use a real-world dataset—baseball performance statistics—throughout the book Work with examples of several analytic patterns, and learn when and where you might use them

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

Get Free Practical Hadoop By Example

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*

Special Features:

- Introduction to MapReduce
- Examples illustrating ideas in practice
- Hadoop's Streaming API
- Other related tools, like Pig and Hive

About The Book: *Hadoop in Action* introduces the subject and teaches you how to write programs in the MapReduce style. It starts with a few easy examples and then moves quickly to show Hadoop use in more complex data analysis tasks. Included are best practices and design patterns of MapReduce programming. This book requires basic Java skills. Knowing basic statistical concepts can help with the more advanced examples.

Get a jump start on using Azure HDInsight and

Get Free Practical Hadoop By Example

Hadoop Ecosystem components. As most Hadoop and Big Data projects are written in either Java, Scala, or Python, this book minimizes the effort to learn another language and is written from the perspective of a .NET developer. Hadoop components are covered, including Hive, Pig, HBase, Storm, and Spark on Azure HDInsight, and code samples are written in .NET only. Processing Big Data with Azure HDInsight covers the fundamentals of big data, how businesses are using it to their advantage, and how Azure HDInsight fits into the big data world. This book introduces Hadoop and big data concepts and then dives into creating different solutions with HDInsight and the Hadoop Ecosystem. It covers concepts with real-world scenarios and code examples, making sure you get hands-on experience. The best way to utilize this book is to practice while reading. After reading this book you will be familiar with Azure HDInsight and how it can be utilized to build big data solutions, including batch processing, stream analytics, interactive processing, and storing and retrieving data in an efficient manner. What You'll Learn Understand the fundamentals of HDInsight and Hadoop Work with HDInsight cluster Query with Apache Hive and Apache Pig Store and retrieve data with Apache HBase Stream data processing using Apache Storm Work with Apache Spark Who This Book Is For Software developers, technical

Get Free Practical Hadoop By Example

architects, data scientists/analyts, and Hadoop administrators who want to develop on Microsoft's managed Hadoop offering, HDInsight

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

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

Get Free Practical Hadoop By Example

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.

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

Get Free Practical Hadoop By Example

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.

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

Practical Hadoop Security is an excellent resource for administrators planning a production Hadoop

Get Free Practical Hadoop By Example

deployment who want to secure their Hadoop clusters. A detailed guide to the security options and configuration within Hadoop itself, author Bhushan Lakhe takes you through a comprehensive study of how to implement defined security within a Hadoop cluster in a hands-on way. You will start with a detailed overview of all the security options available for Hadoop, including popular extensions like Kerberos and OpenSSH, and then delve into a hands-on implementation of user security (with illustrated code samples) with both in-the-box features and with security extensions implemented by leading vendors. No security system is complete without a monitoring and tracing facility, so Practical Hadoop Security next steps you through audit logging and monitoring technologies for Hadoop, as well as ready to use implementation and configuration examples--again with illustrated code samples. The book concludes with the most important aspect of Hadoop security – encryption. Both types of encryptions, for data in transit and data at rest, are discussed at length with leading open source projects that integrate directly with Hadoop at no licensing cost. Practical Hadoop Security: Explains importance of security, auditing and encryption within a Hadoop installation Describes how the leading players have incorporated these features within their Hadoop distributions and provided extensions Demonstrates how to set up and use these features to your benefit and make your Hadoop installation secure without impacting performance or ease of use

"Apache Hadoop is helping drive the Big Data revolution. Now, its data processing has been completely

Get Free Practical Hadoop By Example

overhauled: Apache Hadoop YARN provides resource management at data center scale and easier ways to create distributed applications that process petabytes of data. And now in Apache Hadoop™ YARN, two Hadoop technical leaders show you how to develop new applications and adapt existing code to fully leverage these revolutionary advances." -- From the Amazon Re-architect relational applications to NoSQL, integrate relational database management systems with the Hadoop ecosystem, and transform and migrate relational data to and from Hadoop components. This book covers the best-practice design approaches to re-architecting your relational applications and transforming your relational data to optimize concurrency, security, denormalization, and performance. Winner of IBM's 2012 Gerstner Award for his implementation of big data and data warehouse initiatives and author of Practical Hadoop Security, author Bhushan Lakhe walks you through the entire transition process. First, he lays out the criteria for deciding what blend of re-architecting, migration, and integration between RDBMS and HDFS best meets your transition objectives. Then he demonstrates how to design your transition model. Lakhe proceeds to cover the selection criteria for ETL tools, the implementation steps for migration with SQOOP- and Flume-based data transfers, and transition optimization techniques for tuning partitions, scheduling aggregations, and redesigning ETL. Finally, he assesses the pros and cons of data lakes and Lambda architecture as integrative solutions and illustrates their implementation with real-world case studies.

Get Free Practical Hadoop By Example

Hadoop/NoSQL solutions do not offer by default certain relational technology features such as role-based access control, locking for concurrent updates, and various tools for measuring and enhancing performance. Practical Hadoop Migration shows how to use open-source tools to emulate such relational functionalities in Hadoop ecosystem components. What You'll Learn Decide whether you should migrate your relational applications to big data technologies or integrate them Transition your relational applications to Hadoop/NoSQL platforms in terms of logical design and physical implementation Discover RDBMS-to-HDFS integration, data transformation, and optimization techniques Consider when to use Lambda architecture and data lake solutions Select and implement Hadoop-based components and applications to speed transition, optimize integrated performance, and emulate relational functionalities Who This Book Is For Database developers, database administrators, enterprise architects, Hadoop/NoSQL developers, and IT leaders. Its secondary readership is project and program managers and advanced students of database and management information systems. Let Hadoop For Dummies help harness the power of your data and rein in the information overload Big data has become big business, and companies and organizations of all sizes are struggling to find ways to retrieve valuable information from their massive data sets with becoming overwhelmed. Enter Hadoop and this easy-to-understand For Dummies guide. Hadoop For Dummies helps readers understand the value of big data, make a business case for using Hadoop, navigate

Get Free Practical Hadoop By Example

the Hadoop ecosystem, and build and manage Hadoop applications and clusters. Explains the origins of Hadoop, its economic benefits, and its functionality and practical applications Helps you find your way around the Hadoop ecosystem, program MapReduce, utilize design patterns, and get your Hadoop cluster up and running quickly and easily Details how to use Hadoop applications for data mining, web analytics and personalization, large-scale text processing, data science, and problem-solving Shows you how to improve the value of your Hadoop cluster, maximize your investment in Hadoop, and avoid common pitfalls when building your Hadoop cluster From programmers challenged with building and maintaining affordable, scaleable data systems to administrators who must deal with huge volumes of information effectively and efficiently, this how-to has something to help you with Hadoop.

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

The Complete Guide to Data Science with Hadoop—For Technical Professionals, Businesspeople, and Students Demand is soaring for professionals who can solve real data science problems with Hadoop and Spark. Practical Data Science with Hadoop® and Spark is your complete guide to doing just that. Drawing on immense experience with Hadoop and big data, three leading experts bring together everything you need: high-level concepts, deep-dive techniques, real-world use

Get Free Practical Hadoop By Example

cases, practical applications, and hands-on tutorials. The authors introduce the essentials of data science and the modern Hadoop ecosystem, explaining how Hadoop and Spark have evolved into an effective platform for solving data science problems at scale. In addition to comprehensive application coverage, the authors also provide useful guidance on the important steps of data ingestion, data munging, and visualization. Once the groundwork is in place, the authors focus on specific applications, including machine learning, predictive modeling for sentiment analysis, clustering for document analysis, anomaly detection, and natural language processing (NLP). This guide provides a strong technical foundation for those who want to do practical data science, and also presents business-driven guidance on how to apply Hadoop and Spark to optimize ROI of data science initiatives. Learn What data science is, how it has evolved, and how to plan a data science career How data volume, variety, and velocity shape data science use cases Hadoop and its ecosystem, including HDFS, MapReduce, YARN, and Spark Data importation with Hive and Spark Data quality, preprocessing, preparation, and modeling Visualization: surfacing insights from huge data sets Machine learning: classification, regression, clustering, and anomaly detection Algorithms and Hadoop tools for predictive modeling Cluster analysis and similarity functions Large-scale

Get Free Practical Hadoop By Example

anomaly detection NLP: applying data science to human language

The professional's one-stop guide to this open-source, Java-based big data framework Professional Hadoop is the complete reference and resource for experienced developers looking to employ Apache Hadoop in real-world settings. Written by an expert team of certified Hadoop developers, committers, and Summit speakers, this book details every key aspect of Hadoop technology to enable optimal processing of large data sets. Designed expressly for the professional developer, this book skips over the basics of database development to get you acquainted with the framework's processes and capabilities right away. The discussion covers each key Hadoop component individually, culminating in a sample application that brings all of the pieces together to illustrate the cooperation and interplay that make Hadoop a major big data solution.

Coverage includes everything from storage and security to computing and user experience, with expert guidance on integrating other software and more. Hadoop is quickly reaching significant market usage, and more and more developers are being called upon to develop big data solutions using the Hadoop framework. This book covers the process from beginning to end, providing a crash course for professionals needing to learn and apply Hadoop quickly. Configure storage, UE, and in-memory

Get Free Practical Hadoop By Example

computing Integrate Hadoop with other programs including Kafka and Storm Master the fundamentals of Apache Big Top and Ignite Build robust data security with expert tips and advice Hadoop's popularity is largely due to its accessibility. Open-source and written in Java, the framework offers almost no barrier to entry for experienced database developers already familiar with the skills and requirements real-world programming entails. Professional Hadoop gives you the practical information and framework-specific skills you need quickly.

This book is a complete practical approach for Hadoop lovers. It is mainly aimed at beginners who want to have a hands-on experience with Hadoop and its ecosystem. Its simplicity and step-by-step explanation will help students and other readers in the computer science industry to use this book as a reference manual. The book has been divided into various chapters that cover Hadoop installation, Summary on Hadoop core components, General commands in Hadoop with examples, SQOOP-import & export commands with verification steps, Pig Latin Commands, Analysis using Pig Latin, Pig Script examples, HiveQL Queries and expected outputs and HBase with CRUD operations. In short, this book is a guide for programmers and non-programmers to begin their projects in Hadoop. It is also suitable as a reference manual for students and

Get Free Practical Hadoop By Example

professionals who are new to the Hadoop Ecosystems.

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.

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

Get Free Practical Hadoop By Example

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

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

Get Free Practical Hadoop By Example

stories Use basic tools and techniques to handle backup and catastrophic failure

Filled with practical, step-by-step instructions and clear explanations for the most important and useful tasks. This is a Packt Instant How-to guide, which provides concise and clear recipes for getting started with Hadoop. This book is for big data enthusiasts and would-be Hadoop programmers. It is also meant for Java programmers who either have not worked with Hadoop at all, or who know Hadoop and MapReduce but are not sure how to deepen their understanding.

As more corporations turn to Hadoop to store and process their most valuable data, the risk of a potential breach of those systems increases exponentially. This practical book not only shows Hadoop administrators and security architects how to protect Hadoop data from unauthorized access, it also shows how to limit the ability of an attacker to corrupt or modify data in the event of a security breach. Authors Ben Spivey and Joey Echeverria provide in-depth information about the security features available in Hadoop, and organize them according to common computer security concepts. You'll also get real-world examples that demonstrate how you can apply these concepts to your use cases. Understand the challenges of securing distributed systems, particularly Hadoop Use best practices for preparing Hadoop cluster

Get Free Practical Hadoop By Example

hardware as securely as possible Get an overview of the Kerberos network authentication protocol Delve into authorization and accounting principles as they apply to Hadoop Learn how to use mechanisms to protect data in a Hadoop cluster, both in transit and at rest Integrate Hadoop data ingest into enterprise-wide security architecture Ensure that security architecture reaches all the way to end-user access This book is a practical guide on using 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 is a practical tutorial on using a 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 these all work together as a cohesive big data development platform. What you'll learn How to set up environment in Linux for Hadoop projects using Cloudera Hadoop Distribution CDH 5. How to run a MapReduce job How to store data with Apache Hive, Apache HBase How to index data in HDFS with Apache Solr How to develop a Kafka messaging system How to develop a Mahout User Recommender System How to stream Logs to HDFS with Apache Flume How to transfer data from MySQL database to Hive, HDFS and HBase with Sqoop How create a Hive table over Apache Solr Who this book is for: The primary audience is Apache Hadoop developers. Pre-requisite knowledge of Linux and some knowledge of Hadoop is required. Perform forensic investigations on Hadoop clusters with cutting-edge tools and techniques About This Book Identify,

Get Free Practical Hadoop By Example

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

Get Free Practical Hadoop By Example

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.

This IBM® Redbooks® publication provides topics to help the technical community take advantage of the resilience, scalability, and performance of the IBM Power Systems™ platform to implement or integrate an IBM Data Engine for Hadoop and Spark solution for analytics solutions to access, manage, and analyze data sets to improve business outcomes. This book documents topics to demonstrate and take advantage of the analytics strengths of the IBM POWER8® platform, the IBM analytics software portfolio, and selected third-party tools to help solve customer's data analytic workload requirements. This book describes how to plan, prepare, install, integrate, manage, and show how to use the IBM Data Engine for Hadoop and Spark solution to run analytic workloads on IBM POWER8. In addition, this publication delivers documentation to complement available IBM analytics solutions to help your data analytic needs. This

Get Free Practical Hadoop By Example

publication strengthens the position of IBM analytics and big data solutions with a well-defined and documented deployment model within an IBM POWER8 virtualized environment so that customers have a planned foundation for security, scaling, capacity, resilience, and optimization for analytics workloads. This book is targeted at technical professionals (analytics consultants, technical support staff, IT Architects, and IT Specialists) that are responsible for delivering analytics solutions and support on IBM Power Systems.

Presents information and techniques of using Hadoop to query and analyze data which is distributed across large clusters.

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

Get Free Practical Hadoop By Example

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.

Understand data in a simple way using a data lake. **KEY FEATURES** ? In-depth practical demonstration of Hadoop/Yarn concepts with numerous examples. ? Includes graphical illustrations and visual explanations for Hadoop commands and parameters. ? Includes details of dimensional modeling and Data Vault modeling. ? Includes details of how to create and define a structure to a data lake.

DESCRIPTION The book 'Data Processing and Modeling with Hadoop' explains how a distributed system works and its benefits in the big data era in a straightforward and clear manner. After reading the book, you will be able to plan and organize projects involving a massive amount of data. The book describes the standards and technologies that aid in data management and compares them to other technology business standards. The reader receives practical guidance on how to segregate and separate data into zones, as well as how to develop a model that can aid in data evolution. It discusses security and the measures that are utilized to reduce the impact of security. Self-service analytics, Data Lake, Data Vault 2.0, and Data Mesh are discussed in the book. After reading this book, the reader will have a thorough understanding of how to structure a data lake, as well as the ability to plan, organize, and carry out the implementation of a

Get Free Practical Hadoop By Example

data-driven business with full governance and security.

WHAT YOU WILL LEARN ? Learn the basics of components to the Hadoop Ecosystem. ? Understand the structure, files, and zones of a Data Lake. ? Learn to implement the security part of the Hadoop Ecosystem. ? Learn to work with the Data Vault 2.0 modeling. ? Learn to develop a strategy to define good governance. ? Learn new tools to work with Data and Big Data WHO THIS BOOK IS FOR This book caters to big data developers, technical specialists, consultants, and students who want to build good proficiency in big data.

Knowing basic SQL concepts, modeling, and development would be good, although not mandatory. TABLE OF CONTENTS 1. Understanding the Current Moment 2. Defining the Zones 3. The Importance of Modeling 4. Massive Parallel Processing 5. Doing ETL/ELT 6. A Little Governance 7. Talking About Security 8. What Are the Next Steps? Learn practical uses for some of the hottest tech applications trending among technology professionals We are living in an era of digital revolution. On the horizon, many emerging digital technologies are being developed at a breathtaking speed. Whether we like it or not, whether we are ready or not, digital technologies are going to penetrate more and more, deeper and deeper, into every aspect of our lives. This is going to fundamentally change how we live, how we work, and how we socialize. Java, as a modern high-level programming language, is an excellent tool for helping us to learn these digital technologies, as well as to develop digital applications, such as IoT, AI, Cybersecurity, Blockchain and more. Practical Java Programming uses Java as a tool to help you learn these new digital technologies and to be better prepared for the future changes. Gives you a brief overview for getting started with Java Programming Dives into how you can apply your new knowledge to some of the biggest trending applications today Helps you understand how to

Get Free Practical Hadoop By Example

program Java to interact with operating systems, networking, and mobile applications Shows you how Java can be used in trending tech applications such as IoT (Internet of Things), AI (Artificial Intelligence), Cybersecurity, and Blockchain Get ready to find out firsthand how Java can be used for connected home devices, healthcare, the cloud, and all the hottest tech applications.

This edited book first consolidates the results of the EU-funded EDISON project (Education for Data Intensive Science to Open New science frontiers), which developed training material and information to assist educators, trainers, employers, and research infrastructure managers in identifying, recruiting and inspiring the data science professionals of the future. It then deepens the presentation of the information and knowledge gained to allow for easier assimilation by the reader. The contributed chapters are presented in sequence, each chapter picking up from the end point of the previous one. After the initial book and project overview, the chapters present the relevant data science competencies and body of knowledge, the model curriculum required to teach the required foundations, profiles of professionals in this domain, and use cases and applications. The text is supported with appendices on related process models. The book can be used to develop new courses in data science, evaluate existing modules and courses, draft job descriptions, and plan and design efficient data-intensive research teams across scientific disciplines.

[Copyright: da5d9eb8497004b61b32d804531f3aba](https://www.edisondatascience.eu/)