

## Hbase Administration Cookbook

Deployment Guide for InfoSphere GuardiumHadoop MapReduce CookbookAzure Serverless Computing CookbookIBM DB2 9.7 Advanced Administration CookbookMicrosoft Azure Essentials - Fundamentals of AzureApache Sqoop CookbookIBM PowerHA SystemMirror for AIX CookbookSQL Server 2016 Reporting Services CookbookDatabase AdministrationZooKeeperData-intensive Text Processing with MapReduceHadoop OperationsHadoop Operations and Cluster Management CookbookData Science from ScratchBig Data For DummiesHBase in ActionInfrastructure as CodeData Lake Development with Big DataLearning HBaseExpert Hadoop 2 AdministrationElectronics CookbookBig Data Analytics with JavaMastering MongoDB 3.xHBase EssentialsBig Data AnalyticsBriggsHadoop 2.x Administration CookbookHBase High Performance CookbookHadoop Real-World Solutions CookbookMongoDB and PythonMastering Apache SparkCloudera Administration HandbookGetting Started with KuduHadoop: The Definitive GuideBig Data Analytics with R and HadoopApache Spark 2.x Machine Learning CookbookHbase Administration CookbookGetting Started with ImpalaHBase: The Definitive GuideLearning YARN

## Deployment Guide for InfoSphere Guardium

Virtualization, cloud, containers, server automation, and software-defined networking are meant to simplify IT operations. But many organizations adopting these technologies have found that it only leads to a faster-growing sprawl of unmanageable systems. This is where infrastructure as code can help. With this practical guide, author Kief Morris of ThoughtWorks shows you how to effectively use principles, practices, and patterns pioneered through the DevOps movement to manage cloud age infrastructure. Ideal for system administrators, infrastructure engineers, team leads, and architects, this book demonstrates various tools, techniques, and patterns you can use to implement infrastructure as code. In three parts, you'll learn about the platforms and tooling involved in creating and configuring infrastructure elements, patterns for using these tools, and practices for making infrastructure as code work in your environment. Examine the pitfalls that organizations fall into when adopting the new generation of infrastructure technologies Understand the capabilities and service models of dynamic infrastructure platforms Learn about tools that provide, provision, and configure core infrastructure resources Explore services and tools for managing a dynamic infrastructure Learn specific patterns and practices for provisioning servers, building server templates, and updating running servers

## Hadoop MapReduce Cookbook

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.

## **Azure Serverless Computing Cookbook**

Learn the basics of analytics on big data using Java, machine learning and other big data tools About This Book Acquire real-world set of tools for building enterprise level data science applications Surpasses the barrier of other languages in data science and learn create useful object-oriented codes Extensive use of Java compliant big data tools like apache spark, Hadoop, etc. Who This Book Is For This book is for Java developers who are looking to perform data analysis in production environment. Those who wish to implement data analysis in their Big data applications will find this book helpful. What You Will Learn Start from simple analytic tasks on big data Get into more complex tasks with predictive analytics on big data using machine learning Learn real time analytic tasks Understand the concepts with examples and case studies Prepare and refine data for analysis Create charts in order to understand the data See various real-world datasets In Detail This book covers case studies such as sentiment analysis on a tweet dataset, recommendations on a movielens dataset, customer segmentation on an ecommerce dataset, and graph analysis on actual flights dataset. This book is an end-to-end guide to implement analytics on big data with Java. Java is the de facto language for major big data environments, including Hadoop. This book will teach you how to perform analytics on big data with production-friendly Java. This book basically divided into two sections. The first part is an introduction that will help the readers get acquainted with big data environments, whereas the second part will contain a hardcore discussion on all the concepts in analytics on big data. It will take you from data analysis and data visualization to the core concepts and advantages of machine learning, real-life usage of regression and classification using Naive Bayes, a deep discussion on the concepts of clustering, and a review of simple neural networks on big data using deepLearning4j or plain Java Spark code. This book is a must-have book for Java developers who want to start learning big data analytics and want to use it in the real world. Style and approach The approach of book is to deliver practical learning modules in manageable content. Each chapter is a self-contained unit of a concept in big data analytics. Book will step by step builds the competency in the area of big data analytics. Examples using real world case studies to give ideas of real applications and how to use the techniques mentioned. The examples and case studies will be shown using both theory and code.

## **IBM DB2 9.7 Advanced Administration Cookbook**

This is a practical hands-on book with clear instructions and lot of code examples. It takes a simple approach, guiding you through different architectural topics using realistic sample projects

## **Microsoft Azure Essentials - Fundamentals of Azure**

This book is intended for developers and Big Data engineers who want to know all about HBase at a hands-on level. For in-depth understanding, it would be helpful to have a bit of familiarity with HDFS and MapReduce programming concepts with no prior experience with HBase or similar technologies. This book is also for Big Data enthusiasts and database developers who have worked with other NoSQL

databases and now want to explore HBase as another futuristic, scalable database solution in the Big Data space.

### **Apache Sqoop Cookbook**

Over 50 recipes to help you build applications hosted on Serverless architecture using Azure Functions. About This Book Enhance Azure Functions with continuous deployment using Visual Studio Team Services Learn to deploy and manage cost-effective and highly available serverless applications using Azure Functions This recipe-based guide will teach you to build a robust serverless environment Who This Book Is For If you are a Cloud administrator, architect, or developer who wants to build scalable systems and deploy serverless applications with Azure functions, then this book is for you. Prior knowledge and hands-on experience with core services of Microsoft Azure is required. What You Will Learn Develop different event-based handlers supported by serverless architecture supported by Microsoft Cloud Platform - Azure Integrate Azure Functions with different Azure Services to develop Enterprise-level applications Get to know the best practices in organizing and refactoring the code within the Azure functions Test, troubleshoot, and monitor the Azure functions to deliver high-quality, reliable, and robust cloud-centric applications Automate mundane tasks at various levels right from development to deployment and maintenance Learn how to develop statefulserverless applications and also self-healing jobs using DurableFunctions In Detail Microsoft provides a solution to easily run small segment of code in the Cloud with Azure Functions. Azure Functions provides solutions for processing data, integrating systems, and building simple APIs and microservices. The book starts with intermediate-level recipes on serverless computing along with some use cases on benefits and key features of Azure Functions. Then, we'll deep dive into the core aspects of Azure Functions such as the services it provides, how you can develop and write Azure functions, and how to monitor and troubleshoot them. Moving on, you'll get practical recipes on integrating DevOps with Azure functions, and providing continuous integration and continuous deployment with Visual Studio Team Services. It also provides hands-on steps and tutorials based on real-world serverless use cases, to guide you through configuring and setting up your serverless environments with ease. Finally, you'll see how to manage Azure functions, providing enterprise-level security and compliance to your serverless code architecture. By the end of this book, you will have all the skills required to work with serverless code architecture, providing continuous delivery to your users. Style and approach This recipe-based guide explains the different features of Azure Function by taking a real-world application related to a specific domain. You will learn how to implement automation and DevOps and discover industry best practices to develop applications hosted on serverless architecture using Azure functions.

### **IBM PowerHA SystemMirror for AIX Cookbook**

Our world is being revolutionized by data-driven methods: access to large amounts of data has generated new insights and opened exciting new opportunities in commerce, science, and computing applications. Processing the enormous quantities of data necessary for these advances requires large clusters, making distributed computing paradigms more crucial than ever. MapReduce is a

programming model for expressing distributed computations on massive datasets and an execution framework for large-scale data processing on clusters of commodity servers. The programming model provides an easy-to-understand abstraction for designing scalable algorithms, while the execution framework transparently handles many system-level details, ranging from scheduling to synchronization to fault tolerance. This book focuses on MapReduce algorithm design, with an emphasis on text processing algorithms common in natural language processing, information retrieval, and machine learning. We introduce the notion of MapReduce design patterns, which represent general reusable solutions to commonly occurring problems across a variety of problem domains. This book not only intends to help the reader "think in MapReduce", but also discusses limitations of the programming model as well. This volume is a printed version of a work that appears in the Synthesis Digital Library of Engineering and Computer Science. Synthesis Lectures provide concise, original presentations of important research and development topics, published quickly, in digital and print formats. For more information visit [www.morganclaypool.com](http://www.morganclaypool.com)

### **SQL Server 2016 Reporting Services Cookbook**

Exciting projects that will teach you how complex data can be exploited to gain maximum insights About This Book Architect a good HBase cluster for a very large distributed system Get to grips with the concepts of performance tuning with HBase A practical guide full of engaging recipes and attractive screenshots to enhance your system's performance Who This Book Is For This book is intended for developers and architects who want to know all about HBase at a hands-on level. This book is also for big data enthusiasts and database developers who have worked with other NoSQL databases and now want to explore HBase as another futuristic scalable database solution in the big data space. What You Will Learn Configure HBase from a high performance perspective Grab data from various RDBMS/Flat files into the HBASE systems Understand table design and perform CRUD operations Find out how the communication between the client and server happens in HBase Grasp when to use and avoid MapReduce and how to perform various tasks with it Get to know the concepts of scaling with HBase through practical examples Set up Hbase in the Cloud for a small scale environment Integrate HBase with other tools including ElasticSearch In Detail Apache HBase is a non-relational NoSQL database management system that runs on top of HDFS. It is an open source, distributed, versioned, column-oriented store and is written in Java to provide random real-time access to big Data. We'll start off by ensuring you have a solid understanding the basics of HBase, followed by giving you a thorough explanation of architecting a HBase cluster as per our project specifications. Next, we will explore the scalable structure of tables and we will be able to communicate with the HBase client. After this, we'll show you the intricacies of MapReduce and the art of performance tuning with HBase. Following this, we'll explain the concepts pertaining to scaling with HBase. Finally, you will get an understanding of how to integrate HBase with other tools such as ElasticSearch. By the end of this book, you will have learned enough to exploit HBase for boost system performance. Style and approach This book is intended for software quality assurance/testing professionals, software project managers, or software developers with prior experience in using Selenium and Java to test web-based applications. This books also provides examples for C#, Python, and Ruby users.

## Database Administration

Individual self-contained code recipes. Solve specific problems using individual recipes, or work through the book to develop your capabilities. If you are a big data enthusiast and striving to use Hadoop to solve your problems, this book is for you. Aimed at Java programmers with some knowledge of Hadoop MapReduce, this is also a comprehensive reference for developers and system admins who want to get up to speed using Hadoop.

## ZooKeeper

Data science libraries, frameworks, modules, and toolkits are great for doing data science, but they're also a good way to dive into the discipline without actually understanding data science. In this book, you'll learn how many of the most fundamental data science tools and algorithms work by implementing them from scratch. If you have an aptitude for mathematics and some programming skills, author Joel Grus will help you get comfortable with the math and statistics at the core of data science, and with hacking skills you need to get started as a data scientist. Today's messy glut of data holds answers to questions no one's even thought to ask. This book provides you with the know-how to dig those answers out. Get a crash course in Python Learn the basics of linear algebra, statistics, and probability—and understand how and when they're used in data science Collect, explore, clean, munge, and manipulate data Dive into the fundamentals of machine learning Implement models such as k-nearest Neighbors, Naive Bayes, linear and logistic regression, decision trees, neural networks, and clustering Explore recommender systems, natural language processing, network analysis, MapReduce, and databases

## Data-intensive Text Processing with MapReduce

Fast data ingestion, serving, and analytics in the Hadoop ecosystem have forced developers and architects to choose solutions using the least common denominator—either fast analytics at the cost of slow data ingestion or fast data ingestion at the cost of slow analytics. There is an answer to this problem. With the Apache Kudu column-oriented data store, you can easily perform fast analytics on fast data. This practical guide shows you how. Begun as an internal project at Cloudera, Kudu is an open source solution compatible with many data processing frameworks in the Hadoop environment. In this book, current and former solutions professionals from Cloudera provide use cases, examples, best practices, and sample code to help you get up to speed with Kudu. Explore Kudu's high-level design, including how it spreads data across servers Fully administer a Kudu cluster, enable security, and add or remove nodes Learn Kudu's client-side APIs, including how to integrate Apache Impala, Spark, and other frameworks for data manipulation Examine Kudu's schema design, including basic concepts and primitives necessary to make your project successful Explore case studies for using Kudu for real-time IoT analytics, predictive modeling, and in combination with another storage engine

## Hadoop Operations

Explore architectural approaches to building Data Lakes that ingest, index, manage, and analyze massive amounts of data using Big Data technologies About This Book Comprehend the intricacies of architecting a Data Lake and build a data strategy around your current data architecture Efficiently manage vast amounts of data and deliver it to multiple applications and systems with a high degree of performance and scalability Packed with industry best practices and use-case scenarios to get you up-and-running Who This Book Is For This book is for architects and senior managers who are responsible for building a strategy around their current data architecture, helping them identify the need for a Data Lake implementation in an enterprise context. The reader will need a good knowledge of master data management and information lifecycle management, and experience of Big Data technologies. What You Will Learn Identify the need for a Data Lake in your enterprise context and learn to architect a Data Lake Learn to build various tiers of a Data Lake, such as data intake, management, consumption, and governance, with a focus on practical implementation scenarios Find out the key considerations to be taken into account while building each tier of the Data Lake Understand Hadoop-oriented data transfer mechanism to ingest data in batch, micro-batch, and real-time modes Explore various data integration needs and learn how to perform data enrichment and data transformations using Big Data technologies Enable data discovery on the Data Lake to allow users to discover the data Discover how data is packaged and provisioned for consumption Comprehend the importance of including data governance disciplines while building a Data Lake In Detail A Data Lake is a highly scalable platform for storing huge volumes of multistructured data from disparate sources with centralized data management services. This book explores the potential of Data Lakes and explores architectural approaches to building data lakes that ingest, index, manage, and analyze massive amounts of data using batch and real-time processing frameworks. It guides you on how to go about building a Data Lake that is managed by Hadoop and accessed as required by other Big Data applications. This book will guide readers (using best practices) in developing Data Lake's capabilities. It will focus on architect data governance, security, data quality, data lineage tracking, metadata management, and semantic data tagging. By the end of this book, you will have a good understanding of building a Data Lake for Big Data. Style and approach Data Lake Development with Big Data provides architectural approaches to building a Data Lake. It follows a use case-based approach where practical implementation scenarios of each key component are explained. It also helps you understand how these use cases are implemented in a Data Lake. The chapters are organized in a way that mimics the sequential data flow evidenced in a Data Lake.

## **Hadoop Operations and Cluster Management Cookbook**

Over 100 practical recipes to help you become an expert Hadoop administrator About This Book Become an expert Hadoop administrator and perform tasks to optimize your Hadoop Cluster Import and export data into Hive and use Oozie to manage workflow. Practical recipes will help you plan and secure your Hadoop cluster, and make it highly available Who This Book Is For If you are a system administrator with a basic understanding of Hadoop and you want to get into Hadoop administration, this book is for you. It's also ideal if you are a Hadoop administrator who wants a quick reference guide to all the Hadoop administration-related tasks and solutions to commonly occurring problems What You Will Learn

Set up the Hadoop architecture to run a Hadoop cluster smoothly Maintain a Hadoop cluster on HDFS, YARN, and MapReduce Understand high availability with Zookeeper and Journal Node Configure Flume for data ingestion and Oozie to run various workflows Tune the Hadoop cluster for optimal performance Schedule jobs on a Hadoop cluster using the Fair and Capacity scheduler Secure your cluster and troubleshoot it for various common pain points In Detail Hadoop enables the distributed storage and processing of large datasets across clusters of computers. Learning how to administer Hadoop is crucial to exploit its unique features. With this book, you will be able to overcome common problems encountered in Hadoop administration. The book begins with laying the foundation by showing you the steps needed to set up a Hadoop cluster and its various nodes. You will get a better understanding of how to maintain Hadoop cluster, especially on the HDFS layer and using YARN and MapReduce. Further on, you will explore durability and high availability of a Hadoop cluster. You'll get a better understanding of the schedulers in Hadoop and how to configure and use them for your tasks. You will also get hands-on experience with the backup and recovery options and the performance tuning aspects of Hadoop. Finally, you will get a better understanding of troubleshooting, diagnostics, and best practices in Hadoop administration. By the end of this book, you will have a proper understanding of working with Hadoop clusters and will also be able to secure, encrypt it, and configure auditing for your Hadoop clusters. Style and approach This book contains short recipes that will help you run a Hadoop cluster efficiently. The recipes are solutions to real-life problems that administrators encounter while working with a Hadoop cluster

### **Data Science from Scratch**

Microsoft Azure Essentials from Microsoft Press is a series of free ebooks designed to help you advance your technical skills with Microsoft Azure. The first ebook in the series, Microsoft Azure Essentials: Fundamentals of Azure, introduces developers and IT professionals to the wide range of capabilities in Azure. The authors - both Microsoft MVPs in Azure - present both conceptual and how-to content for key areas, including: Azure Websites and Azure Cloud Services Azure Virtual Machines Azure Storage Azure Virtual Networks Databases Azure Active Directory Management tools Business scenarios Watch Microsoft Press's blog and Twitter (@MicrosoftPress) to learn about other free ebooks in the "Microsoft Azure Essentials" series.

### **Big Data For Dummies**

Building distributed applications is difficult enough without having to coordinate the actions that make them work. This practical guide shows how Apache ZooKeeper helps you manage distributed systems, so you can focus mainly on application logic. Even with ZooKeeper, implementing coordination tasks is not trivial, but this book provides good practices to give you a head start, and points out caveats that developers and administrators alike need to watch for along the way. In three separate sections, ZooKeeper contributors Flavio Junqueira and Benjamin Reed introduce the principles of distributed systems, provide ZooKeeper programming techniques, and include the information you need to administer this service. Learn how ZooKeeper solves common coordination tasks Explore the ZooKeeper API's Java and C implementations and how they differ Use methods to

track and react to ZooKeeper state changes Handle failures of the network, application processes, and ZooKeeper itself Learn about ZooKeeper's trickier aspects dealing with concurrency, ordering, and configuration Use the Curator high-level interface for connection management Become familiar with ZooKeeper internals and administration tools

### **HBase in Action**

A handy reference guide for data analysts and data scientists to help to obtain value from big data analytics using Spark on Hadoop clusters About This Book This book is based on the latest 2.0 version of Apache Spark and 2.7 version of Hadoop integrated with most commonly used tools. Learn all Spark stack components including latest topics such as DataFrames, DataSets, GraphFrames, Structured Streaming, DataFrame based ML Pipelines and SparkR. Integrations with frameworks such as HDFS, YARN and tools such as Jupyter, Zeppelin, NiFi, Mahout, HBase Spark Connector, GraphFrames, H2O and Hivemall. Who This Book Is For Though this book is primarily aimed at data analysts and data scientists, it will also help architects, programmers, and practitioners. Knowledge of either Spark or Hadoop would be beneficial. It is assumed that you have basic programming background in Scala, Python, SQL, or R programming with basic Linux experience. Working experience within big data environments is not mandatory. What You Will Learn Find out and implement the tools and techniques of big data analytics using Spark on Hadoop clusters with wide variety of tools used with Spark and Hadoop Understand all the Hadoop and Spark ecosystem components Get to know all the Spark components: Spark Core, Spark SQL, DataFrames, DataSets, Conventional and Structured Streaming, MLLib, ML Pipelines and Graphx See batch and real-time data analytics using Spark Core, Spark SQL, and Conventional and Structured Streaming Get to grips with data science and machine learning using MLLib, ML Pipelines, H2O, Hivemall, Graphx, SparkR and Hivemall. In Detail Big Data Analytics book aims at providing the fundamentals of Apache Spark and Hadoop. All Spark components – Spark Core, Spark SQL, DataFrames, Data sets, Conventional Streaming, Structured Streaming, MLLib, Graphx and Hadoop core components – HDFS, MapReduce and Yarn are explored in greater depth with implementation examples on Spark + Hadoop clusters. It is moving away from MapReduce to Spark. So, advantages of Spark over MapReduce are explained at great depth to reap benefits of in-memory speeds. DataFrames API, Data Sources API and new Data set API are explained for building Big Data analytical applications. Real-time data analytics using Spark Streaming with Apache Kafka and HBase is covered to help building streaming applications. New Structured streaming concept is explained with an IOT (Internet of Things) use case. Machine learning techniques are covered using MLLib, ML Pipelines and SparkR and Graph Analytics are covered with GraphX and GraphFrames components of Spark. Readers will also get an opportunity to get started with web based notebooks such as Jupyter, Apache Zeppelin and data flow tool Apache NiFi to analyze and visualize data. Style and approach This step-by-step pragmatic guide will make life easy no matter what your level of experience. You will deep dive into Apache Spark on Hadoop clusters through ample exciting real-life examples. Practical tutorial explains data science in simple terms to help programmers and data analysts get started with Data Science

## Infrastructure as Code

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

## Data Lake Development with Big Data

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

## Learning HBase

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

### **Expert Hadoop 2 Administration**

A thorough reference on database administration outlines a variety of DBA roles and responsibilities and discusses such topics as data modeling and normalization, database/application design, change management, database security and data integrity, performance issues, disaster planning, and other essentials. Original. (Advanced)

### **Electronics Cookbook**

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

### **Big Data Analytics with Java**

If you're looking for a scalable storage solution to accommodate a virtually endless amount of data, this book shows you how Apache HBase can fulfill your needs. As the open source implementation of Google's BigTable architecture, HBase scales to billions of rows and millions of columns, while ensuring that write and read performance remain constant. Many IT executives are asking pointed questions about HBase. This book provides meaningful answers, whether you're evaluating this non-relational database or planning to put it into practice right away. Discover how tight integration with Hadoop makes scalability with HBase easier Distribute large datasets across an inexpensive cluster of commodity servers Access HBase with native Java clients, or with gateway servers providing REST, Avro, or Thrift APIs Get details on HBase's architecture, including the storage format, write-ahead log, background processes, and more Integrate HBase with Hadoop's MapReduce framework for massively parallelized data processing jobs Learn how to tune clusters, design schemas, copy tables, import bulk data, decommission nodes, and many other tasks

### **Mastering MongoDB 3.x**

Gain expertise in processing and storing data by using advanced techniques with Apache Spark About This Book Explore the integration of Apache Spark with third party applications such as H2O, Databricks and Titan Evaluate how Cassandra and Hbase can be used for storage An advanced guide with a combination of instructions and practical examples to extend the most up-to date Spark functionalities Who This Book Is For If you are a developer with some experience with Spark and want to strengthen your knowledge of how to get around in the world of Spark, then this book is ideal for you. Basic knowledge of Linux, Hadoop and Spark is assumed. Reasonable knowledge of Scala is expected. What You Will Learn Extend the tools available for processing and storage Examine clustering and classification using MLlib Discover Spark stream processing via Flume, HDFS Create a schema in Spark SQL, and learn how a Spark schema can be populated with data Study Spark based graph processing using Spark GraphX Combine Spark with H2O and deep learning and learn why it is useful Evaluate how graph storage works with Apache Spark, Titan, HBase and Cassandra Use Apache Spark in the cloud with Databricks and AWS In Detail Apache Spark is an in-memory cluster based parallel processing system that provides a wide range of functionality like graph processing, machine learning, stream processing and SQL. It operates at unprecedented speeds, is easy to use and offers a rich set of data transformations. This book aims to take your limited knowledge of Spark to the next level by teaching you how to expand Spark functionality. The book commences with an overview of the Spark eco-system. You will learn how to use MLlib to create a fully working neural net for handwriting recognition. You will then discover how stream processing can be tuned for optimal performance and to ensure parallel processing. The book extends to show how to incorporate H2O for machine learning, Titan for graph based storage, Databricks for cloud-based Spark. Intermediate Scala based code examples are provided for Apache Spark module processing in a CentOS Linux and Databricks cloud environment. Style and approach This book is an extensive guide to Apache Spark modules and tools and shows how Spark's functionality can be extended for real-time processing and storage with worked examples.

## HBase Essentials

If you're among the many hobbyists and designers who came to electronics through Arduino and Raspberry Pi, this cookbook will help you learn and apply the basics of electrical engineering without the need for an EE degree. Through a series of practical recipes, you'll learn how to solve specific problems while diving into as much or as little theory as you're comfortable with. Author Simon Monk (Raspberry Pi Cookbook) breaks down this complex subject into several topics, from using the right transistor to building and testing projects and prototypes. With this book, you can quickly search electronics topics and go straight to the recipe you need. It also serves as an ideal reference for experienced electronics makers. This cookbook includes: Theoretical concepts such as Ohm's law and the relationship between power, voltage, and current The fundamental use of resistors, capacitors and inductors, diodes, transistors and integrated circuits, and switches and relays Recipes on power, sensors and motors, integrated circuits, and radio frequency for designing electronic circuits and devices Advice on using Arduino and Raspberry Pi in electronics projects How to build and use tools, including multimeters, oscilloscopes, simulations software, and unsoldered prototypes

## Big Data Analytics

IBM® InfoSphere® Guardium® provides the simplest, most robust solution for data security and data privacy by assuring the integrity of trusted information in your data center. InfoSphere Guardium helps you reduce support costs by automating the entire compliance auditing process across heterogeneous environments. InfoSphere Guardium offers a flexible and scalable solution to support varying customer architecture requirements. This IBM Redbooks® publication provides a guide for deploying the Guardium solutions. This book also provides a roadmap process for implementing an InfoSphere Guardium solution that is based on years of experience and best practices that were collected from various Guardium experts. We describe planning, installation, configuration, monitoring, and administering an InfoSphere Guardium environment. We also describe use cases and how InfoSphere Guardium integrates with other IBM products. The guidance can help you successfully deploy and manage an IBM InfoSphere Guardium system. This book is intended for the system administrators and support staff who are responsible for deploying or supporting an InfoSphere Guardium environment.

## Briggs

Learn how to write, tune, and port SQL queries and other statements for a Big Data environment, using Impala—the massively parallel processing SQL query engine for Apache Hadoop. The best practices in this practical guide help you design database schemas that not only interoperate with other Hadoop components, and are convenient for administrators to manage and monitor, but also accommodate future expansion in data size and evolution of software capabilities. Written by John Russell, documentation lead for the Cloudera Impala project, this book gets you working with the most recent Impala releases quickly. Ideal for database developers and business analysts, the latest revision covers analytics functions, complex types, incremental statistics, subqueries, and submission to the Apache incubator. Getting Started with Impala includes advice from Cloudera's development team, as well as insights from its consulting engagements with customers. Learn how Impala integrates with a wide range of Hadoop components. Attain high performance and scalability for huge data sets on production clusters. Explore common developer tasks, such as porting code to Impala and optimizing performance. Use tutorials for working with billion-row tables, date- and time-based values, and other techniques. Learn how to transition from rigid schemas to a flexible model that evolves as needs change. Take a deep dive into joins and the roles of statistics.

## Hadoop 2.x Administration Cookbook

Solve specific problems using individual self-contained code recipes, or work through the book to develop your capabilities. This book is packed with easy-to-follow code and commands used for illustration, which makes your learning curve easy and quick. If you are a Hadoop cluster system administrator with Unix/Linux system management experience and you are looking to get a good grounding in how to set up and manage a Hadoop cluster, then this book is for you. It's assumed

that you will have some experience in Unix/Linux command line already, as well as being familiar with network communication basics.

### **HBase High Performance Cookbook**

This IBM® Redbooks® publication can help you install, tailor, and configure the new IBM PowerHA® Version 7.1.3, and understand new and improved features such as migrations, cluster administration, and advanced topics like configuring in a virtualized environment including workload partitions (WPARs). With this book, you can gain a broad understanding of the IBM PowerHA SystemMirror® architecture. If you plan to install, migrate, or administer a high availability cluster, this book is right for you. This book can help IBM AIX® professionals who seek a comprehensive and task-oriented guide for developing the knowledge and skills required for PowerHA cluster design, implementation, and daily system administration. It provides a combination of theory and practical experience. This book is targeted toward technical professionals (consultants, technical support staff, IT architects, and IT specialists) who are responsible for providing high availability solutions and support with the IBM PowerHA SystemMirror Standard on IBM POWER® systems.

### **Hadoop Real-World Solutions Cookbook**

An easy-to-follow Apache Hadoop administrator's guide filled with practical screenshots and explanations for each step and configuration. This book is great for administrators interested in setting up and managing a large Hadoop cluster. If you are an administrator, or want to be an administrator, and you are ready to build and maintain a production-level cluster running CDH5, then this book is for you.

### **MongoDB and Python**

As part of Packt's cookbook series, each recipe offers a practical, step-by-step solution to common problems found in HBase administration. This book is for HBase administrators, developers, and will even help Hadoop administrators. You are not required to have HBase experience, but are expected to have a basic understanding of Hadoop and MapReduce.

### **Mastering Apache Spark**

Find the right big data solution for your business or organization Big data management is one of the major challenges facing business, industry, and not-for-profit organizations. Data sets such as customer transactions for a mega-retailer, weather patterns monitored by meteorologists, or social network activity can quickly outpace the capacity of traditional data management tools. If you need to develop or manage big data solutions, you'll appreciate how these four experts define, explain, and guide you through this new and often confusing concept. You'll learn what it is, why it matters, and how to choose and implement solutions that work. Effectively managing big data is an issue of growing importance to businesses, not-for-profit organizations, government, and IT professionals Authors

are experts in information management, big data, and a variety of solutions. Explains big data in detail and discusses how to select and implement a solution, security concerns to consider, data storage and presentation issues, analytics, and much more. Provides essential information in a no-nonsense, easy-to-understand style that is empowering. Big Data For Dummies cuts through the confusion and helps you take charge of big data solutions for your organization.

### **Cloudera Administration Handbook**

Moving beyond MapReduce - learn resource management and big data processing using YARN. About This Book: Deep dive into YARN components, schedulers, life cycle management and security architecture. Create your own Hadoop-YARN applications and integrate big data technologies with YARN. Step-by-step guide to provision, manage, and monitor Hadoop-YARN clusters with ease. Who This Book Is For: This book is intended for those who want to understand what YARN is and how to efficiently use it for the resource management of large clusters. For cluster administrators, this book gives a detailed explanation of provisioning and managing YARN clusters. If you are a Java developer or an open source contributor, this book will help you to drill down the YARN architecture, write your own YARN applications and understand the application execution phases. This book will also help big data engineers explore YARN integration with real-time analytics technologies such as Spark and Storm. What You Will Learn: Explore YARN features and offerings. Manage big data clusters efficiently using the YARN framework. Create single as well as multi-node Hadoop-YARN clusters on Linux machines. Understand YARN components and their administration. Gain insights into application execution flow over a YARN cluster. Write your own distributed application and execute it over YARN cluster. Work with schedulers and queues for efficient scheduling of applications. Integrate big data projects like Spark and Storm with YARN. In Detail: Today enterprises generate huge volumes of data. In order to provide effective services and to make smarter and more intelligent decisions from these huge volumes of data, enterprises use big-data analytics. In recent years, Hadoop has been used for massive data storage and efficient distributed processing of data. The Yet Another Resource Negotiator (YARN) framework solves the design problems related to resource management faced by the Hadoop 1.x framework by providing a more scalable, efficient, flexible, and highly available resource management framework for distributed data processing. This book starts with an overview of the YARN features and explains how YARN provides a business solution for growing big data needs. You will learn to provision and manage single, as well as multi-node, Hadoop-YARN clusters in the easiest way. You will walk through the YARN administration, life cycle management, application execution, REST APIs, schedulers, security framework and so on. You will gain insights about the YARN components and features such as ResourceManager, NodeManager, ApplicationMaster, Container, Timeline Server, High Availability, Resource Localisation and so on. The book explains Hadoop-YARN commands and the configurations of components and explores topics such as High Availability, Resource Localization and Log aggregation. You will then be ready to develop your own ApplicationMaster and execute it over a Hadoop-YARN cluster. Towards the end of the book, you will learn about the security architecture and integration of YARN with big data technologies like Spark and Storm. This book promises conceptual as well as practical knowledge of resource management using YARN.

Style and approach Starting with the basics and covering the core concepts with the practical usage, this tutorial is a complete guide to learn and explore YARN offerings.

### **Getting Started with Kudu**

Summary HBase in Action has all the knowledge you need to design, build, and run applications using HBase. First, it introduces you to the fundamentals of distributed systems and large scale data handling. Then, you'll explore real-world applications and code samples with just enough theory to understand the practical techniques. You'll see how to build applications with HBase and take advantage of the MapReduce processing framework. And along the way you'll learn patterns and best practices. About the Technology HBase is a NoSQL storage system designed for fast, random access to large volumes of data. It runs on commodity hardware and scales smoothly from modest datasets to billions of rows and millions of columns. About this Book HBase in Action is an experience-driven guide that shows you how to design, build, and run applications using HBase. First, it introduces you to the fundamentals of handling big data. Then, you'll explore HBase with the help of real applications and code samples and with just enough theory to back up the practical techniques. You'll take advantage of the MapReduce processing framework and benefit from seeing HBase best practices in action. 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. What's Inside When and how to use HBase Practical examples Design patterns for scalable data systems Deployment, integration, and design Written for developers and architects familiar with data storage and processing. No prior knowledge of HBase, Hadoop, or MapReduce is required. Table of Contents PART 1 HBASE FUNDAMENTALS Introducing HBase Getting started Distributed HBase, HDFS, and MapReduce PART 2 ADVANCED CONCEPTS HBase table design Extending HBase with coprocessors Alternative HBase clients PART 3 EXAMPLE APPLICATIONS HBase by example: OpenTSDB Scaling GIS on HBase PART 4 OPERATIONALIZING HBASE Deploying HBase Operations

### **Hadoop: The Definitive Guide**

If you are an administrator or developer who wants to enter the world of Big Data and BigTables and would like to learn about HBase, this is the book for you.

### **Big Data Analytics with R and Hadoop**

An expert's guide to build fault tolerant MongoDB application About This Book Master the advanced modeling, querying, and administration techniques in MongoDB and become a MongoDB expert Covers the latest updates and Big Data features frequently used by professional MongoDB developers and administrators If your goal is to become a certified MongoDB professional, this book is your perfect companion Who This Book Is For Mastering MongoDB is a book for database developers, architects, and administrators who want to learn how to use MongoDB more effectively and productively. If you have experience in, and are interested in working with, NoSQL databases to build apps and websites, then this book is for

you. What You Will Learn Get hands-on with advanced querying techniques such as indexing, expressions, arrays, and more. Configure, monitor, and maintain highly scalable MongoDB environment like an expert. Master replication and data sharding to optimize read/write performance. Design secure and robust applications based on MongoDB. Administer MongoDB-based applications on-premise or in the cloud Scale MongoDB to achieve your design goals Integrate MongoDB with big data sources to process huge amounts of data In Detail MongoDB has grown to become the de facto NoSQL database with millions of users—from small startups to Fortune 500 companies. Addressing the limitations of SQL schema-based databases, MongoDB pioneered a shift of focus for DevOps and offered sharding and replication maintainable by DevOps teams. The book is based on MongoDB 3.x and covers topics ranging from database querying using the shell, built in drivers, and popular ODM mappers to more advanced topics such as sharding, high availability, and integration with big data sources. You will get an overview of MongoDB and how to play to its strengths, with relevant use cases. After that, you will learn how to query MongoDB effectively and make use of indexes as much as possible. The next part deals with the administration of MongoDB installations on-premise or in the cloud. We deal with database internals in the next section, explaining storage systems and how they can affect performance. The last section of this book deals with replication and MongoDB scaling, along with integration with heterogeneous data sources. By the end this book, you will be equipped with all the required industry skills and knowledge to become a certified MongoDB developer and administrator. Style and approach This book takes a practical, step-by-step approach to explain the concepts of MongoDB. Practical use-cases involving real-world examples are used throughout the book to clearly explain theoretical concepts.

## **Apache Spark 2.x Machine Learning Cookbook**

Create interactive cross-platform reports and dashboards using SQL Server 2016 Reporting Services About This Book Get up to speed with the newly-introduced enhancements and the more advanced query and reporting features Easily access your important data by creating visually appealing dashboards in the Power BI practical recipe Create cross-browser and cross-platform reports using SQL Server 2016 Reporting Services Who This Book Is For This book is for software professionals who develop and implement reporting solutions using Microsoft SQL Server. It is especially relevant for professionals who are software engineers, software architects, DW/BI engineers, and DW/BI architects who perform simple to complex report authoring implementations. This book is also suitable for those who develop software solutions that integrate reporting solutions and are keen to learn about Microsoft SQL Server 2016's features and capabilities. What You Will Learn Key capabilities, architecture, and components of Reporting Services New features that have been added to Reporting Services Design the architecture for reporting solutions Design the architecture for BI solutions Implement reporting solutions using Reporting Services Improve the performance, availability, and scalability of the reporting solution Enhance reporting solutions with custom programming and improved security In Detail Microsoft SQL Server 2016 Reporting Services comes with many new features. It offers different types of reporting such as Production, Ad-hoc, Dashboard, Mash-up, and Analytical. SQL Server 2016 also has a surfeit of new features including Mobile Reporting, and Power BI integration. This book

contains recipes that explore the new and advanced features added to SQL Server 2016. The first few chapters cover recipes on configuring components and how to explore these new features. You'll learn to build your own reporting solution with data tools and report builder, along with learning techniques to create visually appealing reports. This book also has recipes for enhanced mobile reporting solutions, accessing these solutions effectively, and delivering interactive business intelligence solutions. Towards the end of the book, you'll get to grips with running reporting services in SharePoint integrated mode and be able to administer, monitor, and secure your reporting solution. This book covers about the new offerings of Microsoft SQL Server 2016 Reporting Services in comprehensive detail and uses examples of real-world problem-solving business scenarios. Style and approach This comprehensive cookbook follows a problem-solution approach to help you overcome any obstacle when creating interactive, visually-appealing reports using SQL Server 2016 Reporting Services. Each recipe focuses on a specific task and is written in a clear, solution-focused style.

### **Hbase Administration Cookbook**

Simplify machine learning model implementations with Spark About This Book Solve the day-to-day problems of data science with Spark This unique cookbook consists of exciting and intuitive numerical recipes Optimize your work by acquiring, cleaning, analyzing, predicting, and visualizing your data Who This Book Is For This book is for Scala developers with a fairly good exposure to and understanding of machine learning techniques, but lack practical implementations with Spark. A solid knowledge of machine learning algorithms is assumed, as well as hands-on experience of implementing ML algorithms with Scala. However, you do not need to be acquainted with the Spark ML libraries and ecosystem. What You Will Learn Get to know how Scala and Spark go hand-in-hand for developers when developing ML systems with Spark Build a recommendation engine that scales with Spark Find out how to build unsupervised clustering systems to classify data in Spark Build machine learning systems with the Decision Tree and Ensemble models in Spark Deal with the curse of high-dimensionality in big data using Spark Implement Text analytics for Search Engines in Spark Streaming Machine Learning System implementation using Spark In Detail Machine learning aims to extract knowledge from data, relying on fundamental concepts in computer science, statistics, probability, and optimization. Learning about algorithms enables a wide range of applications, from everyday tasks such as product recommendations and spam filtering to cutting edge applications such as self-driving cars and personalized medicine. You will gain hands-on experience of applying these principles using Apache Spark, a resilient cluster computing system well suited for large-scale machine learning tasks. This book begins with a quick overview of setting up the necessary IDEs to facilitate the execution of code examples that will be covered in various chapters. It also highlights some key issues developers face while working with machine learning algorithms on the Spark platform. We progress by uncovering the various Spark APIs and the implementation of ML algorithms with developing classification systems, recommendation engines, text analytics, clustering, and learning systems. Toward the final chapters, we'll focus on building high-end applications and explain various unsupervised methodologies and challenges to tackle when implementing with big data ML systems. Style and approach This book is packed with intuitive recipes supported with line-by-line

explanations to help you understand how to optimize your work flow and resolve problems when working with complex data modeling tasks and predictive algorithms. This is a valuable resource for data scientists and those working on large scale data projects.

### **Getting Started with Impala**

Hadoop: The Definitive Guide helps you harness the power of your data. Ideal for processing large datasets, the Apache Hadoop framework is an open source implementation of the MapReduce algorithm on which Google built its empire. This comprehensive resource demonstrates how to use Hadoop to build reliable, scalable, distributed systems: programmers will find details for analyzing large datasets, and administrators will learn how to set up and run Hadoop clusters. Complete with case studies that illustrate how Hadoop solves specific problems, this book helps you: Use the Hadoop Distributed File System (HDFS) for storing large datasets, and run distributed computations over those datasets using MapReduce Become familiar with Hadoop's data and I/O building blocks for compression, data integrity, serialization, 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 Use Pig, a high-level query language for large-scale data processing Take advantage of HBase, Hadoop's database for structured and semi-structured data Learn ZooKeeper, a toolkit of coordination primitives for building distributed systems If you have lots of data -- whether it's gigabytes or petabytes -- Hadoop is the perfect solution. Hadoop: The Definitive Guide is the most thorough book available on the subject. "Now you have the opportunity to learn about Hadoop from a master-not only of the technology, but also of common sense and plain talk."-- Doug Cutting, Hadoop Founder, Yahoo!

### **HBase: The Definitive Guide**

How do you start? How should you build a plan for cloud migration for your entire portfolio? How will your organization be affected by these changes? This book, based on real-world cloud experiences by enterprise IT teams, seeks to provide the answers to these questions. Here, you'll see what makes the cloud so compelling to enterprises; with which applications you should start your cloud journey; how your organization will change, and how skill sets will evolve; how to measure progress; how to think about security, compliance, and business buy-in; and how to exploit the ever-growing feature set that the cloud offers to gain strategic and competitive advantage.

### **Learning YARN**

Learn how to leverage MongoDB with your Python applications, using the hands-on recipes in this book. You get complete code samples for tasks such as making fast geo queries for location-based apps, efficiently indexing your user documents for social-graph lookups, and many other scenarios. This guide explains the basics of the document-oriented database and shows you how to set up a Python environment with it. Learn how to read and write to MongoDB, apply idiomatic

MongoDB and Python patterns, and use the database with several popular Python web frameworks. You'll discover how to model your data, write effective queries, and avoid concurrency problems such as race conditions and deadlocks. The recipes will help you: Read, write, count, and sort documents in a MongoDB collection Learn how to use the rich MongoDB query language Maintain data integrity in replicated/distributed MongoDB environments Use embedding to efficiently model your data without joins Code defensively to avoid keyerrors and other bugs Apply atomic operations to update game scores, billing systems, and more with the fast accounting pattern Use MongoDB with the Pylons 1.x, Django, and Pyramid web frameworks

[ROMANCE](#) [ACTION & ADVENTURE](#) [MYSTERY & THRILLER](#) [BIOGRAPHIES & HISTORY](#) [CHILDREN'S](#) [YOUNG ADULT](#) [FANTASY](#) [HISTORICAL FICTION](#) [HORROR](#) [LITERARY FICTION](#) [NON-FICTION](#) [SCIENCE FICTION](#)