

HEADFIRST HADOOP EDITION (DOWNLOAD ONLY)

Hadoop: The Definitive Guide

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

Hadoop in Practice

Summary Hadoop in Practice, Second Edition provides over 100 tested, instantly useful techniques that will help you conquer big data, using Hadoop. This revised new edition covers changes and new features in the Hadoop core architecture, including MapReduce 2. Brand new chapters cover YARN and integrating Kafka, Impala, and Spark SQL with Hadoop. You'll also get new and updated techniques for Flume, Sqoop, and Mahout, all of which have seen major new versions recently. In short, this is the most practical, up-to-date coverage of Hadoop available anywhere. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book It's always a good time to upgrade your Hadoop skills! Hadoop in Practice, Second Edition provides a collection of 104 tested, instantly useful techniques for analyzing real-time streams, moving data securely, machine learning, managing large-scale clusters, and taming big data using Hadoop. This completely revised edition covers changes and new features in Hadoop core, including MapReduce 2 and YARN. You'll pick up hands-on best practices for integrating Spark, Kafka, and Impala with Hadoop, and get new and updated techniques for the latest versions of Flume, Sqoop, and Mahout. In short, this is the most practical, up-to-date coverage of Hadoop available. Readers need to know a programming language like Java and have basic familiarity with Hadoop. What's Inside Thoroughly updated for Hadoop 2 How to write YARN applications Integrate real-time technologies like Storm, Impala, and Spark Predictive analytics using Mahout and RR Readers need to know a programming language like Java and have basic familiarity with Hadoop. About the Author Alex Holmes works on tough big-data problems. He is a software engineer, author, speaker, and blogger specializing in large-scale Hadoop projects. Table of Contents PART 1 BACKGROUND AND FUNDAMENTALS Hadoop in a heartbeat Introduction to YARN PART 2 DATA LOGISTICS Data serialization—working with text and beyond Organizing and optimizing data in HDFS Moving data into and out of Hadoop PART 3 BIG DATA PATTERNS Applying MapReduce patterns to big data Utilizing data structures and algorithms at scale Tuning, debugging, and testing PART 4 BEYOND MAPREDUCE SQL on Hadoop Writing a YARN application

Hadoop MapReduce v2 Cookbook - Second Edition

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

Apache Hadoop 3 Quick Start Guide

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

Hadoop: The Definitive Guide

Get ready to unlock the power of your data. With the fourth edition of 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. Using Hadoop 2 exclusively, author Tom White presents new chapters on YARN and several Hadoop-related projects such as Parquet, Flume, Crunch, and Spark. You'll learn about recent changes to Hadoop, and explore new case studies on Hadoop's role in healthcare systems and genomics data processing. Learn fundamental components such as MapReduce, HDFS, and YARN Explore MapReduce in depth, including steps for developing applications with it Set up and maintain a Hadoop cluster running HDFS and MapReduce on YARN Learn two data formats: Avro for data serialization and Parquet for nested data Use data ingestion tools such as Flume (for streaming data) and Sqoop (for bulk data transfer) Understand how high-level data processing tools like Pig, Hive, Crunch, and Spark work with Hadoop Learn the HBase distributed database and the ZooKeeper distributed configuration service

Hadoop Administrator Interview Questions

Cloudera® Enterprise is one of the fastest growing platforms for the BigData computing world, which accommodate various open source tools like CDH, Hive, Impala, HBase and many more as well as licensed

products like Cloudera Manager and Cloudera Navigator. There are various organization who had already deployed the Cloudera Enterprise solution in the production env, and running millions of queries and data processing on daily basis. Cloudera Enterprise is such a vast and managed platform, that as individual, cannot manage the entire cluster. Even single administrator cannot have entire cluster knowledge, that's the reason there is a huge demand for the Cloudera Administrator in the market specially in the North America, Canada, France, UAE, Germany, India etc. Many international investment and retail bank already installed the Cloudera Enterprise in the production environment, Healthcare and retail e-commerce industry which has huge volume of data generated on daily basis do not have a choice and they have to have Hadoop based platform deployed. Cloudera Enterprise is the pioneer and not any other company is close to the Cloudera for the Hadoop Solution, and demand for Cloudera certified Hadoop Administrators are high in demand. That's the reason HadoopExam is launching Hadoop Administrator Interview Preparation Material, which is specially designed for the Cloudera Enterprise product, you have to go through all the questions mentioned in this book before your real interview. This book certainly helpful for your real interview, however does not guarantee that you will clear that interview or not. In this book we have covered various terminology, concepts, architectural perspective, Impala, Hive, Cloudera Manager, Cloudera Navigator and Some part of Cloudera Altus. We will be continuously upgrading this book. So, you can get the access to most recent material. Please keep in mind this book is written mainly for the Cloudera Enterprise Hadoop Administrator, and it may be helpful if you are working on any other Hadoop Solution provider as well.

Hadoop 2 Quick-Start Guide

Get Started Fast with Apache Hadoop® 2, YARN, and Today's Hadoop Ecosystem With Hadoop 2.x and YARN, Hadoop moves beyond MapReduce to become practical for virtually any type of data processing. Hadoop 2.x and the Data Lake concept represent a radical shift away from conventional approaches to data usage and storage. Hadoop 2.x installations offer unmatched scalability and breakthrough extensibility that supports new and existing Big Data analytics processing methods and models. Hadoop® 2 Quick-Start Guide is the first easy, accessible guide to Apache Hadoop 2.x, YARN, and the modern Hadoop ecosystem. Building on his unsurpassed experience teaching Hadoop and Big Data, author Douglas Eadline covers all the basics you need to know to install and use Hadoop 2 on personal computers or servers, and to navigate the powerful technologies that complement it. Eadline concisely introduces and explains every key Hadoop 2 concept, tool, and service, illustrating each with a simple "beginning-to-end" example and identifying trustworthy, up-to-date resources for learning more. This guide is ideal if you want to learn about Hadoop 2 without getting mired in technical details. Douglas Eadline will bring you up to speed quickly, whether you're a user, admin, devops specialist, programmer, architect, analyst, or data scientist. Coverage Includes Understanding what Hadoop 2 and YARN do, and how they improve on Hadoop 1 with MapReduce Understanding Hadoop-based Data Lakes versus RDBMS Data Warehouses Installing Hadoop 2 and core services on Linux machines, virtualized sandboxes, or clusters Exploring the Hadoop Distributed File System (HDFS) Understanding the essentials of MapReduce and YARN application programming Simplifying programming and data movement with Apache Pig, Hive, Sqoop, Flume, Oozie, and HBase Observing application progress, controlling jobs, and managing workflows Managing Hadoop efficiently with Apache Ambari—including recipes for HDFS to NFSv3 gateway, HDFS snapshots, and YARN configuration Learning basic Hadoop 2 troubleshooting, and installing Apache Hue and Apache Spark

Apache Flume

Summary The Spark distributed data processing platform provides an easy-to-implement tool for ingesting, streaming, and processing data from any source. In Spark in Action, Second Edition, you'll learn to take advantage of Spark's core features and incredible processing speed, with applications including real-time computation, delayed evaluation, and machine learning. Spark skills are a hot commodity in enterprises worldwide, and with Spark's powerful and flexible Java APIs, you can reap all the benefits without first learning Scala or Hadoop. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Analyzing enterprise data starts by reading,

filtering, and merging files and streams from many sources. The Spark data processing engine handles this varied volume like a champ, delivering speeds 100 times faster than Hadoop systems. Thanks to SQL support, an intuitive interface, and a straightforward multilanguage API, you can use Spark without learning a complex new ecosystem. About the book *Spark in Action, Second Edition*, teaches you to create end-to-end analytics applications. In this entirely new book, you'll learn from interesting Java-based examples, including a complete data pipeline for processing NASA satellite data. And you'll discover Java, Python, and Scala code samples hosted on GitHub that you can explore and adapt, plus appendixes that give you a cheat sheet for installing tools and understanding Spark-specific terms. What's inside Writing Spark applications in Java Spark application architecture Ingestion through files, databases, streaming, and Elasticsearch Querying distributed datasets with Spark SQL About the reader This book does not assume previous experience with Spark, Scala, or Hadoop. About the author Jean-Georges Perrin is an experienced data and software architect. He is France's first IBM Champion and has been honored for 12 consecutive years. Table of Contents PART 1 - THE THEORY CRIPPLED BY AWESOME EXAMPLES 1 So, what is Spark, anyway? 2 Architecture and flow 3 The majestic role of the dataframe 4 Fundamentally lazy 5 Building a simple app for deployment 6 Deploying your simple app PART 2 - INGESTION 7 Ingestion from files 8 Ingestion from databases 9 Advanced ingestion: finding data sources and building your own 10 Ingestion through structured streaming PART 3 - TRANSFORMING YOUR DATA 11 Working with SQL 12 Transforming your data 13 Transforming entire documents 14 Extending transformations with user-defined functions 15 Aggregating your data PART 4 - GOING FURTHER 16 Cache and checkpoint: Enhancing Spark's performances 17 Exporting data and building full data pipelines 18 Exploring deployment

Spark in Action, Second Edition

Over 90 hands-on recipes to help you learn and master the intricacies of Apache Hadoop 2.X, YARN, Hive, Pig, Oozie, Flume, Sqoop, Apache Spark, and Mahout About This Book- Implement outstanding Machine Learning use cases on your own analytics models and processes.- Solutions to common problems when working with the Hadoop ecosystem.- Step-by-step implementation of end-to-end big data use cases. Who This Book Is For Readers who have a basic knowledge of big data systems and want to advance their knowledge with hands-on recipes. What You Will Learn- Installing and maintaining Hadoop 2.X cluster and its ecosystem.- Write advanced Map Reduce programs and understand design patterns.- Advanced Data Analysis using the Hive, Pig, and Map Reduce programs.- Import and export data from various sources using Sqoop and Flume.- Data storage in various file formats such as Text, Sequential, Parquet, ORC, and RC Files.- Machine learning principles with libraries such as Mahout- Batch and Stream data processing using Apache Spark In Detail Big data is the current requirement. Most organizations produce huge amount of data every day. With the arrival of Hadoop-like tools, it has become easier for everyone to solve big data problems with great efficiency and at minimal cost. Grasping Machine Learning techniques will help you greatly in building predictive models and using this data to make the right decisions for your organization. Hadoop Real World Solutions Cookbook gives readers insights into learning and mastering big data via recipes. The book not only clarifies most big data tools in the market but also provides best practices for using them. The book provides recipes that are based on the latest versions of Apache Hadoop 2.X, YARN, Hive, Pig, Sqoop, Flume, Apache Spark, Mahout and many more such ecosystem tools. This real-world-solution cookbook is packed with handy recipes you can apply to your own everyday issues. Each chapter provides in-depth recipes that can be referenced easily. This book provides detailed practices on the latest technologies such as YARN and Apache Spark. Readers will be able to consider themselves as big data experts on completion of this book. This guide is an invaluable tutorial if you are planning to implement a big data warehouse for your business. Style and approach An easy-to-follow guide that walks you through world of big data. Each tool in the Hadoop ecosystem is explained in detail and the recipes are placed in such a manner that readers can implement them sequentially. Plenty of reference links are provided for advanced reading.

Hadoop Real-World Solutions Cookbook Second Edition

If you want to discover one of the latest tools designed to produce stunning Big Data insights, this book features everything you need to get to grips with your data. Whether you are a data architect, developer, or a business strategist, HDInsight adds value in everything from development, administration, and reporting.

Hdinsight Essentials - Second Edition

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 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?

Data Processing and Modeling with Hadoop

"Get ready to unlock the power of your data. With the fourth edition of 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."

Hadoop

Explore the Hadoop MapReduce v2 ecosystem to gain insights from very large datasets In Detail Starting with installing Hadoop YARN, MapReduce, HDFS, and other Hadoop ecosystem components, with this book, you will soon learn about many exciting topics such as MapReduce patterns, using Hadoop to solve analytics, classifications, online marketing, recommendations, and data indexing and searching. You will learn how to take advantage of Hadoop ecosystem projects including Hive, HBase, Pig, Mahout, Nutch, and Giraph and be introduced to deploying in cloud environments. Finally, you will be able to apply the knowledge you have gained to your own real-world scenarios to achieve the best-possible results. **What You Will Learn** Configure and administer Hadoop YARN, MapReduce v2, and HDFS clusters Use Hive, HBase, Pig, Mahout, and Nutch with Hadoop v2 to solve your big data problems easily and effectively Solve large-scale analytics problems using MapReduce-based applications Tackle complex problems such as classifications, finding relationships, online marketing, recommendations, and searching using Hadoop MapReduce and other related projects Perform massive text data processing using Hadoop MapReduce and other related projects Deploy your clusters to cloud environments Downloading the example code for this book. You can download the example code files for all Packt books you have purchased from your account at

<http://www.PacktPub.com>. If you purchased this book elsewhere, you can visit <http://www.PacktPub.com/support> and register to have the files e-mailed directly to you.

Hadoop MapReduce V2 Cookbook - Second Edition

Get expert guidance on architecting end-to-end data management solutions with Apache Hadoop. While many sources explain how to use various components in the Hadoop ecosystem, this practical book takes you through architectural considerations necessary to tie those components together into a complete tailored application, based on your particular use case. To reinforce those lessons, the book's second section provides detailed examples of architectures used in some of the most commonly found Hadoop applications. Whether you're designing a new Hadoop application, or planning to integrate Hadoop into your existing data infrastructure, Hadoop Application Architectures will skillfully guide you through the process. This book covers: Factors to consider when using Hadoop to store and model data Best practices for moving data in and out of the system Data processing frameworks, including MapReduce, Spark, and Hive Common Hadoop processing patterns, such as removing duplicate records and using windowing analytics Giraph, GraphX, and other tools for large graph processing on Hadoop Using workflow orchestration and scheduling tools such as Apache Oozie Near-real-time stream processing with Apache Storm, Apache Spark Streaming, and Apache Flume Architecture examples for clickstream analysis, fraud detection, and data warehousing

Scaling Big Data with Hadoop and Solr - Second Edition

As data grows exponentially day-by-day, extracting information becomes a tedious activity in itself. Technologies like Hadoop are trying to address some of the concerns, while Solr provides high-speed faceted search. Bringing these two technologies together is helping organizations resolve the problem of information extraction from Big Data by providing excellent distributed faceted search capabilities. Scaling Big Data with Hadoop and Solr is a step-by-step guide that helps you build high performance enterprise search engines while scaling data. Starting with the basics of Apache Hadoop and Solr, this book then dives into advanced topics of optimizing search with some interesting real-world use cases and sample Java code. Scaling Big Data with Hadoop and Solr starts by teaching you the basics of Big Data technologies including Hadoop and its ecosystem and Apache Solr. It explains the different approaches of scaling Big Data with Hadoop and Solr, with discussion regarding the applicability, benefits, and drawbacks of each approach. It then walks readers through how sharding and indexing can be performed on Big Data followed by the performance optimization of Big Data search. Finally, it covers some real-world use cases for Big Data scaling. With this book, you will learn everything you need to know to build a distributed enterprise search platform as well as how to optimize this search to a greater extent resulting in maximum utilization of available resources.

Hadoop Application Architectures

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

Scaling Big Data with Hadoop and Solr

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

Hadoop

Discover how Apache Hadoop can unleash the power of your data. This comprehensive resource shows you how to build and maintain reliable, scalable, distributed systems with the Hadoop framework -- an open source implementation of MapReduce, the algorithm on which Google built its empire. Programmers will find details for analyzing datasets of any size, and administrators will learn how to set up and run Hadoop clusters. This revised edition covers recent changes to Hadoop, including new features such as Hive, Sqoop, and Avro. It also provides illuminating case studies that illustrate how Hadoop is used to solve specific problems. Looking to get the most out of your data? This is your book. Use the Hadoop Distributed File System (HDFS) for storing large datasets, then run distributed computations over those datasets with 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 Analyze datasets with Hive, Hadoop's data warehousing system Take advantage of HBase, Hadoop's database for structured and semi-structured data Learn ZooKeeper, a toolkit of coordination primitives for building distributed systems \"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, Cloudera

Apache Hive Essentials

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research

upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

Hadoop: The Definitive Guide

Learn how to use, deploy, and maintain Apache Spark with this comprehensive guide, written by the creators of the open-source cluster-computing framework. With an emphasis on improvements and new features in Spark 2.0, authors Bill Chambers and Matei Zaharia break down Spark topics into distinct sections, each with unique goals. You'll explore the basic operations and common functions of Spark's structured APIs, as well as Structured Streaming, a new high-level API for building end-to-end streaming applications. Developers and system administrators will learn the fundamentals of monitoring, tuning, and debugging Spark, and explore machine learning techniques and scenarios for employing MLlib, Spark's scalable machine-learning library. Get a gentle overview of big data and Spark Learn about DataFrames, SQL, and Datasets Spark's core APIs through worked examples Dive into Spark's low-level APIs, RDDs, and execution of SQL and DataFrames Understand how Spark runs on a cluster Debug, monitor, and tune Spark clusters and applications Learn the power of Structured Streaming, Spark's stream-processing engine Learn how you can apply MLlib to a variety of problems, including classification or recommendation

Designing Data-Intensive Applications

Tap the power of Big Data with Microsoft technologies Big Data is here, and Microsoft's new Big Data platform is a valuable tool to help your company get the very most out of it. This timely book shows you how to use HDInsight along with HortonWorks Data Platform for Windows to store, manage, analyze, and share Big Data throughout the enterprise. Focusing primarily on Microsoft and HortonWorks technologies but also covering open source tools, Microsoft Big Data Solutions explains best practices, covers on-premises and cloud-based solutions, and features valuable case studies. Best of all, it helps you integrate these new solutions with technologies you already know, such as SQL Server and Hadoop. Walks you through how to integrate Big Data solutions in your company using Microsoft's HDInsight Server, HortonWorks Data Platform for Windows, and open source tools Explores both on-premises and cloud-based solutions Shows how to store, manage, analyze, and share Big Data through the enterprise Covers topics such as Microsoft's approach to Big Data, installing and configuring HortonWorks Data Platform for Windows, integrating Big Data with SQL Server, visualizing data with Microsoft and HortonWorks BI tools, and more Helps you build and execute a Big Data plan Includes contributions from the Microsoft and HortonWorks Big Data product teams If you need a detailed roadmap for designing and implementing a fully deployed Big Data solution, you'll want Microsoft Big Data Solutions.

Spark: The Definitive Guide

Presents an instructional guide to SQL which uses humor and simple images to cover such topics as the structure of relational databases, simple and complex queries, creating multiple tables, and protecting important table data.

Microsoft Big Data Solutions

Learn how to take full advantage of Apache Kafka, the distributed, publish-subscribe queue for handling real-time data feeds. With this comprehensive book, you will understand how Kafka works and how it is designed. Authors Neha Narkhede, Gwen Shapira, and Todd Palino show you how to deploy production Kafka clusters; secure, tune, and monitor them; write rock-solid applications that use Kafka; and build scalable stream-processing applications. Learn how Kafka compares to other queues, and where it fits in the big data ecosystem. Dive into Kafka's internal design Pick up best practices for developing applications that use Kafka. Understand the best way to deploy Kafka in production monitoring, tuning, and maintenance tasks. Learn how to secure a Kafka cluster.

Head First SQL

Now updated for the 2016 PMP exam Learn the latest principles and certification objectives in The PMBOK® Guide, (Fifth Version), in a unique and inspiring way with Head First PMP. This book helps you prepare for the PMP certification exam using a visually rich format designed for the way your brain works. You'll find a full-length sample exam included inside the book. More than just proof of passing a test, a PMP certification means that you have the knowledge to solve most common project problems. But studying for a difficult four-hour exam on project management isn't easy, even for experienced project managers. Drawing on the latest research in neurobiology, cognitive science, and learning theory, Head First PMP offers you a multi-sensory experience that helps the material stick, not a text-heavy approach that puts you to sleep. This book will help you: Learn PMP's underlying concepts to help you understand the PMBOK principles and pass the certification exam with flying colors Get 100% coverage of the latest principles and certification objectives in The PMBOK Guide, Fifth Edition Make use of a thorough and effective preparation guide with hundreds of practice questions and exam strategies Explore the material through puzzles, games, problems, and exercises that make learning easy and entertaining Head First PMP puts project management principles into context to help you understand, remember, and apply them—not just on the exam, but also on the job.

Kafka: The Definitive Guide

Even bad code can function. But if code isn't clean, it can bring a development organization to its knees. Every year, countless hours and significant resources are lost because of poorly written code. But it doesn't have to be that way. Noted software expert Robert C. Martin presents a revolutionary paradigm with Clean Code: A Handbook of Agile Software Craftsmanship. Martin has teamed up with his colleagues from Object Mentor to distill their best agile practice of cleaning code “on the fly” into a book that will instill within you the values of a software craftsman and make you a better programmer—but only if you work at it. What kind of work will you be doing? You'll be reading code—lots of code. And you will be challenged to think about what's right about that code, and what's wrong with it. More importantly, you will be challenged to reassess your professional values and your commitment to your craft. Clean Code is divided into three parts. The first describes the principles, patterns, and practices of writing clean code. The second part consists of several case studies of increasing complexity. Each case study is an exercise in cleaning up code—of transforming a code base that has some problems into one that is sound and efficient. The third part is the payoff: a single chapter containing a list of heuristics and “smells” gathered while creating the case studies. The result is a knowledge base that describes the way we think when we write, read, and clean code. Readers will come away from this book understanding How to tell the difference between good and bad code How to write good code and how to transform bad code into good code How to create good names, good functions, good objects, and good classes How to format code for maximum readability How to implement complete error handling without obscuring code logic How to unit test and practice test-driven development This book is a must for any developer, software engineer, project manager, team lead, or systems analyst with an interest in producing better code.

Head First PMP

Currently used at many colleges, universities, and high schools, this hands-on introduction to computer science is ideal for people with little or no programming experience. The goal of this concise book is not just to teach you Java, but to help you think like a computer scientist. You'll learn how to program—a useful skill by itself—but you'll also discover how to use programming as a means to an end. Authors Allen Downey and Chris Mayfield start with the most basic concepts and gradually move into topics that are more complex, such as recursion and object-oriented programming. Each brief chapter covers the material for one week of a college course and includes exercises to help you practice what you've learned. Learn one concept at a time: tackle complex topics in a series of small steps with examples Understand how to formulate problems, think creatively about solutions, and write programs clearly and accurately Determine which development techniques work best for you, and practice the important skill of debugging Learn relationships among input

and output, decisions and loops, classes and methods, strings and arrays Work on exercises involving word games, graphics, puzzles, and playing cards

Clean Code

What do Ford Financial, IBM, and Victoria's Secret have in common? Enterprise JavaBeans (EJB). As the industry standard for platform-independent reusable business components, EJB has just become Sun Microsystems's latest developer certification. Whether you want to be certifiable or just want to learn the technology inside and out, Head First EJB will get you there in the least painful way. And with the greatest understanding. You'll learn not just what the technology is, but more importantly, why it is, and what it is and isn't good for. You'll learn tricks and tips for EJB development, along with tricks and tips for passing this latest, very challenging Sun Certified Business Component Developer (SCBCD) exam. You'll learn how to think like a server. You'll learn how to think like a bean. And because this is a Head First book, you'll learn how to think about thinking. Co-author Kathy Sierra was one of Sun's first employees to teach brave, early adopter customers how to use EJB. She has the scars. But besides dragging you deep into EJB technology, Kathy and Bert will see you through your certification exam, if you decide to go for it. And nobody knows the certification like they do - they're co-developers of Sun's actual exam! As the second book in the Head First series, Head First EJB follows up the number one best-selling Java book in the US, Head First Java. Find out why reviewers are calling it a revolution in learning tough technical topics, and why Sun Chairman and CEO Scott McNealy says, "Java technology is everywhere...if you develop software and haven't learned Java, it's definitely time to dive in." Head First. And with Head First book, you don't even have to feel guilty about having fun while you're learning; it's all part of the learning theory. If the latest research in cognitive science, education, and neurobiology suggested that boring, dry, and excruciatingly painful was the best way to learn, we'd have done it. Thankfully, it's been shown that your brain has a sense of style, a sense of humour, and a darn good sense of what it likes and dislikes. In Head First EJB, you'll learn all about:

- Component-based and role-based development
- The architecture of EJB, distributed programming with RMI
- Developing and Deploying an EJB application
- The Client View of a Session and Entity bean
- The Session Bean Lifecycle and Component Contract
- The Entity bean Lifecycle and Component Contract
- Container-managed Persistence (CMP)
- Container-managed Relationships (CMR)
- EJB-QL
- Transactions
- Security
- EJB Exceptions
- The Deployment Descriptor
- The Enterprise Bean Environment in JNDI
- Programming Restrictions and Portability

The book includes over 200 mock exam questions that match the tone, style, difficulty, and topics on the real SCBCD exam. See why Kathy and Bert are responsible for thousands of successful exam-passers--"The Sun certification exam was certainly no walk in the park, but Kathy's material allowed me to not only pass the exam, but Ace it!"--Mary Whetsel, Sr. Technology Specialist, Application Strategy and Integration, The St. Paul Companies

"Kathy Sierra and Bert Bates are two of the few people in the world who can make complicated things seem damn simple, and as if that isn't enough, they can make boring things seem interesting."--Paul Wheaton, The Trail Boss, javaranch.com

"Who better to write a Java study guide than Kathy Sierra, reigning queen of Java instruction? Kathy Sierra has done it again. Here is a study guide that almost guarantees you a certification!"--James Cubetta, Systems Engineer, SGI

Think Java

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.

Head First EJB

SQL in a Nutshell applies the eminently useful \"Nutshell\" format to Structured Query Language (SQL), the elegant--but complex--descriptive language that is used to create and manipulate large stores of data. For SQL programmers, analysts, and database administrators, the new second edition of SQL in a Nutshell is the essential date language reference for the world's top SQL database products. SQL in a Nutshell is a lean, focused, and thoroughly comprehensive reference for those who live in a deadline-driven world. This invaluable desktop quick reference drills down and documents every SQL command and how to use it in both commercial (Oracle, DB2, and Microsoft SQL Server) and open source implementations (PostgreSQL, and MySQL). It describes every command and reference and includes the command syntax (by vendor, if the syntax differs across implementations), a clear description, and practical examples that illustrate important concepts and uses. And it also explains how the leading commercial and open sources database product implement SQL. This wealth of information is packed into a succinct, comprehensive, and extraordinarily easy-to-use format that covers the SQL syntax of no less than 4 different databases. When you need fast, accurate, detailed, and up-to-date SQL information, SQL in a Nutshell, Second Edition will be the quick reference you'll reach for every time. SQL in a Nutshell is small enough to keep by your keyboard, and concise (as well as clearly organized) enough that you can look up the syntax you need quickly without having to wade through a lot of useless fluff. You won't want to work on a project involving SQL without it.

Microsoft Azure Essentials - Fundamentals of Azure

This book, written by one of the designers of generics, is a thorough explanation of how to use generics, and particularly, the effect this facility has on the way developers use collections.

SQL in a Nutshell

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

Java Generics and Collections

The software architecture landscape has evolved dramatically over the past decade. Microservices have displaced monoliths. Data and applications are increasingly becoming distributed and decentralised. But composing disparate systems is a hard problem. More recently, software practitioners have been rapidly converging on event-driven architecture as a sustainable way of dealing with complexity - integrating systems without increasing their coupling. In Effective Kafka, Emil Koutanov explores the fundamentals of Event-Driven Architecture - using Apache Kafka - the world's most popular and supported open-source event streaming platform. You'll learn: - The fundamentals of event-driven architecture and event streaming platforms- The background and rationale behind Apache Kafka, its numerous potential uses and applications- The architecture and core concepts - the underlying software components, partitioning and parallelism, load-balancing, record ordering and consistency modes- Installation of Kafka and related tooling - using standalone deployments, clusters, and containerised deployments with Docker- Using CLI tools to interact with and administer Kafka classes, as well as publishing data and browsing topics- Using third-party web-

based tools for monitoring a cluster and gaining insights into the event streams- Building stream processing applications in Java 11 using off-the-shelf client libraries- Patterns and best-practice for organising the application architecture, with emphasis on maintainability and testability of the resulting code- The numerous gotchas that lurk in Kafka's client and broker configuration, and how to counter them- Theoretical background on distributed and concurrent computing, exploring factors affecting their liveness and safety- Best-practices for running multi-tenanted clusters across diverse engineering teams, how teams collaborate to build complex systems at scale and equitably share the cluster with the aid of quotas- Operational aspects of running Kafka clusters at scale, performance tuning and methods for optimising network and storage utilisation- All aspects of Kafka security -including network segregation, encryption, certificates, authentication and authorization. The coverage is progressively delivered and carefully aimed at giving you a journey-like experience into becoming proficient with Apache Kafka and Event-Driven Architecture. The goal is to get you designing and building applications. And by the conclusion of this book, you will be a confident practitioner and a Kafka evangelist within your organisation - wielding the knowledge necessary to teach others.

Hadoop Operations

Data is bigger, arrives faster, and comes in a variety of formats—and it all needs to be processed at scale for analytics or machine learning. But how can you process such varied workloads efficiently? Enter Apache Spark. Updated to include Spark 3.0, this second edition shows data engineers and data scientists why structure and unification in Spark matters. Specifically, this book explains how to perform simple and complex data analytics and employ machine learning algorithms. Through step-by-step walk-throughs, code snippets, and notebooks, you'll be able to: Learn Python, SQL, Scala, or Java high-level Structured APIs Understand Spark operations and SQL Engine Inspect, tune, and debug Spark operations with Spark configurations and Spark UI Connect to data sources: JSON, Parquet, CSV, Avro, ORC, Hive, S3, or Kafka Perform analytics on batch and streaming data using Structured Streaming Build reliable data pipelines with open source Delta Lake and Spark Develop machine learning pipelines with MLlib and productionize models using MLflow

HADOOP

Head First C# is a complete learning experience for learning how to program with C#, XAML, the .NET Framework, and Visual Studio. Fun and highly visual, this introduction to C# is designed to keep you engaged and entertained from first page to last. Updated for Windows 8.1 and Visual Studio 2013, and includes projects for all previous versions of Windows (included in the book, no additional downloading or printing required). You'll build a fully functional video game in the opening chapter, and then learn how to use classes and object-oriented programming, draw graphics and animation, and query data with LINQ and serialize it to files. And you'll do it all by creating games, solving puzzles, and doing hands-on projects. By the time you're done, you'll be a solid C# programmer—and you'll have a great time along the way! Create a fun arcade game in the first chapter, and build games and other projects throughout the book Learn how to use XAML to design attractive and interactive pages and windows Build modern Windows Store apps using the latest Microsoft technology Learn WPF (Windows Presentation Foundation) using the downloadable WPF Learner's Guide Using the Model-View-ViewModel (MVVM) pattern to create robust architecture Build a bonus Windows Phone project and run it in the Visual Studio Windows Phone emulator Projects in the book work with all editions of Visual Studio, including the free Express editions.

Effective Kafka

Java continues to grow and evolve, and this cookbook continues to evolve in tandem. With this guide, you'll get up to speed right away with hundreds of hands-on recipes across a broad range of Java topics. You'll learn useful techniques for everything from string handling and functional programming to network communication. Each recipe includes self-contained code solutions that you can freely use, along with a

discussion of how and why they work. If you're familiar with Java basics, this cookbook will bolster your knowledge of the language and its many recent changes, including how to apply them in your day-to-day development. This updated edition covers changes through Java 12 and parts of 13 and 14. Recipes include: Blade, Laravel's powerful custom templating tool Methods for compiling, running, and debugging Packaging Java classes and building applications Manipulating, comparing, and rearranging text Regular expressions for string and pattern matching Handling numbers, dates, and times Structuring data with collections, arrays, and other types Object-oriented and functional programming techniques Input/output, directory, and filesystem operations Network programming on both client and server Processing JSON for data interchange Multithreading and concurrency Using Java in big data applications Interfacing Java with other languages

Learning Spark

Get a solid grounding in Apache Oozie, the workflow scheduler system for managing Hadoop jobs. With this hands-on guide, two experienced Hadoop practitioners walk you through the intricacies of this powerful and flexible platform, with numerous examples and real-world use cases. Once you set up your Oozie server, you'll dive into techniques for writing and coordinating workflows, and learn how to write complex data pipelines. Advanced topics show you how to handle shared libraries in Oozie, as well as how to implement and manage Oozie's security capabilities. Install and configure an Oozie server, and get an overview of basic concepts Journey through the world of writing and configuring workflows Learn how the Oozie coordinator schedules and executes workflows based on triggers Understand how Oozie manages data dependencies Use Oozie bundles to package several coordinator apps into a data pipeline Learn about security features and shared library management Implement custom extensions and write your own EL functions and actions Debug workflows and manage Oozie's operational details

Head First C#

This book introduces Apache Spark, the open source cluster computing system that makes data analytics fast to write and fast to run. You'll learn how to express parallel jobs with just a few lines of code, and cover applications from simple batch jobs to stream processing and machine learning.--

Java Cookbook

Apache Oozie

[behringer xr 2400 manual](#)

[stihl 110r service manual](#)

[the water footprint assessment manual setting the global standard](#)

[the inkheart trilogy inkspell inkdeath inkworld 1 3 cornelia funke](#)

[2015 core measure pocket guide](#)

[analog digital communication lab manual vtu](#)

[costume since 1945 historical dress from couture to street style](#)

[multistrada 1260 ducati forum](#)

[caribbean women writers essays from the first international conference](#)

[rugerini diesel rd278 manual](#)