
Apache Spark In 24 Hours Sams Teach Yourself

EXPLORE OUR WIDE ARRAY OF APACHE SPARK IN 24 HOURS SAMS TEACH YOURSELF by John D. Sams, Jr. & Kamren

INVITE TO STAGING.EDITOR.SEENIC.IO BOOK SHOP!

At our bookstore, we believe that **Apache Spark In 24 Hours Sams Teach Yourself** have the power to carry us to new worlds and ignite our interests. That's why we provide a **wide array** of **publications** that accommodate all **interests** and **reading degrees**. Whether you're a seasoned visitor or simply beginning your literary journey, we have something for every person.

Our collection consists of both physical and **Apache Spark In 24 Hours Sams Teach Yourself electronic publications**, along with **audiobooks**, so you can choose the layout that ideal suits your choices. We likewise use **individualized publication recommendations** based on your **interests** and past analysis experiences.

Join our area of **publication fans** and **attach** with similar individuals that share your enthusiasm for literary works. We take satisfaction in showcasing brand-new ability and **emerging voices** in the literary world, so you can discover encouraging authors who are pressing limits and redefining categories.

Shop with us today and find the delight and understanding that **Apache Spark In 24 Hours Sams Teach Yourself books** bring.

Trick Takeaways of Apache Spark In 24 Hours Sams Teach Yourself

- Our shop offers a **variety** of Apache Spark In 24 Hours Sams Teach Yourself in numerous layouts, including physical and **electronic books** and **audiobooks**.
- We give **customized publication referrals** based upon your **rate of interests** and previous reading experiences.
- Join our community of **Apache Spark In 24 Hours Sams Teach Yourself book fans** to link with similar people and broaden your literary horizons.
- We showcase brand-new skill and **emerging voices** in the literary world, so you can uncover appealing authors that are pushing limits and redefining styles.
- **Discover** our collection and find the joy and expertise that Apache Spark In 24 Hours Sams Teach Yourself publications bring.

YOURSELF!

At our shop, we pride ourselves on offering a substantial variety of books in numerous genres. Our team believe that everyone needs to have accessibility to literature that talks with them and improves their lives. That's why we make every effort to have something for every person in our collection.

Whether you're a follower of fiction, non-fiction, love, mystery, or any other genre, we have something to please your reading yearnings. Our racks are equipped with both timeless and contemporary titles, and our experienced team is always accessible to help lead you in the right direction.

We believe that analysis is a journey of exploration, and we want to aid you embark on that trip in the very best feasible way. That's why we offer a **wide variety** of Apache Spark In 24 Hours Sams Teach Yourself to choose from, so you can discover the best match for your rate of interests and reviewing preferences.

Explore our wide array of this Apache Spark In 24 Hours Sams Teach Yourself publications!

- Discover a wide range of styles
- Find the **ideal book** for your passions
- Study **exciting narration of Apache Spark In 24 Hours Sams Teach Yourself**
- **Discover brand-new writers** and **emerging voices**
- Locate publications for **every ages** and **analysis levels**

So whether you're seeking a thrilling page-turner or a provocative literary masterpiece, we have something for everyone. Come **explore** our wide variety of publications and let us assist you find your following excellent read!

STUDY FASCINATING STORYTELLING

At our store, our team believe that Apache Spark In 24 Hours Sams Teach Yourself are greater than

SAMS TEACH YOURSELF DIGITAL BOOKS have taken you to new globes and experiences. That's why we carefully curate our option to include books with **fascinating storytelling** that will transport you to different locations and times.

From sprawling legends to heartfelt love, our publications are filled with strong personalities and thought-provoking styles that will certainly keep you involved from start to finish. Whether you're seeking an escape from reality or a deeper understanding of the globe around you, our Apache Spark In 24 Hours Sams Teach Yourself books use something for every person.

Our authors are masters of their craft, using language and images to create brilliant worlds that feel genuine. With every turn of the page, you'll be drawn deeper into the story, anxious to discover what happens following.

Experience the power of exciting storytelling of Apache Spark In 24 Hours Sams Teach Yourself

We believe that analysis is a transformative experience that can transform your life. By diving into a **Apache Spark In 24 Hours Sams Teach Yourself publication** with **fascinating narration**, you have the power to discover brand-new things, get new point of views, and **link** with characters and scenarios that might be various from your own.

Whether you're an experienced reader or simply getting started, our choice of publications with fascinating storytelling makes sure to leave a long lasting impact. We welcome you to **check out** our shop and uncover the magic of storytelling for yourself.

DISCOVER BRAND-NEW WRITERS AND ARISING VOICES

At our shop, we take satisfaction in showcasing brand-new skill and arising voices in the literary world. We believe that every terrific tale deserves to be informed and that the power of literary works depends on its capacity to magnify varied perspectives and voices.

By exploring our choice, you'll have the possibility to discover and support promising authors that are pressing borders and redefining categories. From thought-provoking memoirs to exhilarating debut novels, our collection covers a wide range of styles and styles, guaranteeing that there's something for everyone.

So why not step outdoors your reading convenience zone and uncover a brand-new writer or arising voice today? You never ever understand what literary treasures you might discover!

DELIGHT IN THE CONVENIENCE OF PHYSICAL AND APACHE SPARK IN 24 HOURS

At our shop, we comprehend that every viewers has their own preferences when it comes to the style they select to review their books. That's why we provide both physical and **electronic books of Apache Spark In 24 Hours Sams Teach Yourself** for your **comfort**.

The Comfort of Physical Books

For those who enjoy the feeling of a **Apache Spark In 24 Hours Sams Teach Yourself publication** in their hands and the experience of flipping through the web pages, we offer a wide variety of **physical books**. From hardbounds to books, our choice includes all categories and **reading degrees**. You can surf our shelves and take your time selecting the excellent book to contribute to your collection.

The Convenience of Digital Books

We also understand that some visitors favor the **convenience** of continuing reading a tablet computer or e-reader. That's why we provide a substantial selection of digital books that you can acquire and download and install with ease. You can take your collection with you any place you go and delight in the advantages of having your favorite publications within your reaches.

Whichever style you prefer, we've obtained you covered. Our goal is to give you with the most effective reading experience while likewise keeping in mind your **convenience**. Shop with us today and discover the **delight of reviewing Apache Spark In 24 Hours Sams Teach Yourself** in your recommended layout.

IMPROVE YOUR READING WITH AUDIOBOOKS

At our store, we believe that reading must be a convenient and satisfying experience for everybody. That's why we offer a broad selection of **audiobooks** that you can listen to on-the-go or while unwinding in your home. With our audiobook collection, you can boost your reading experience and submerge yourself in captivating narratives.

Our audiobooks are very carefully curated to offer the very same level of amusement and knowledge as our physical and digital publications. Whether you're paying attention to an awesome mystery or a thought-provoking memoir, our audiobooks will transport you to new globes and provide you a fresh perspective.

The Benefits of Audiobooks

Audiobooks give a special method to enjoy Apache Spark In 24 Hours Sams Teach Yourself and boost your reading experience. Below are some of the advantages of audiobooks:

- Useful for individuals that have trouble checking out due to visual impairment or discovering

impairments.

- Great for multitasking while driving, working out, or doing housework.
- Can improve pronunciation, vocabulary, and language understanding skills.
- Practical for people that deal with holding a publication or analysis for a prolonged time period.

Our audiobook collection includes a range of categories and authors, so you can locate something that straightens with your reading interests. Whether you're a fan of romance, sci-fi, historical fiction, or self-help publications, we have options that satisfy every taste.

Just how to Access Apache Spark In 24 Hours Sams Teach Yourself Audiobook Collection

You can access our audiobook collection via our website or by seeing our physical shop. We offer both physical and digital audiobooks to suit your choices. If you select to purchase Apache Spark In 24 Hours Sams Teach Yourself electronic audiobooks, you can download them promptly and begin listening right away. Alternatively, if you choose physical audiobooks, you can purchase them online and we'll supply them to your doorstep.

Boost your analysis experience with our audiobook collection today and find the pleasure of listening to exciting narrations that will certainly carry you to brand-new worlds.

DISCOVER BOOKS FOR EVERY AGES AND ANALYSIS LEVELS

As a reader, you understand that books can be enjoyed by people of **every ages** and analysis degrees. That is why our collection consists of publications that cater to people of various ages and reading skills.

For the kids, we have a substantial series of youngsters's books, including picture books, storybooks, and activity books. Our selection of young person stories comprises the most recent fads and timeless classics that young people can enjoy.

If you like books for mature readers, we have a wide series of literature that makes sure to stimulate your rate of interest. You can choose from contemporary jobs or delve into classic books from prominent authors.

Whether you are a seasoned bibliophile or just starting on your reading trip, we have a publication for each taste and ability level. Our group of educated personnel can help you discover a Apache Spark In 24 Hours Sams Teach Yourself book that matches your rate of interests and checking out efficiency.

Locate the Perfect Publication for Your Reading Level

At our store, we take pride in dealing with readers of **any ages** and skill degrees. We recognize that choosing the appropriate book can be an overwhelming job, specifically for beginners. That is why we offer individualized suggestions based upon your analysis level and rate of interests.

We want you to feel great in your publication choices, and our team of professionals is here to make that happen. From children's books to adult literature, you can rely on that we will certainly help you discover the ideal publication that matches your reading abilities and rate of interests.

Individualized Recommendations

Our personnel is here to assist you discover your next literary adventure. Connect with us and get individualized suggestions based on your interests and reviewing choices. We'll work with you to locate the excellent book that matches your tastes and leaves you desiring a lot more.

Join our neighborhood of Apache Spark In 24 Hours Sams Teach Yourself publication fans today and connect with similar people who share your enthusiasm for literary works. We can't wait to start brand-new literary journeys with you!

OBTAIN INDIVIDUALIZED BOOK SUGGESTIONS

At our store, we recognize that choosing the ideal book can be overwhelming. That's why we offer customized book suggestions based on your passions and reviewing preferences. Our experienced team is constantly all set to help and provide tailored tips to help you find your following excellent read.

Whether you're seeking a details genre, author, or style, we'll assist you limit your alternatives and find a book that fits your preferences. We put in the time to get to know our consumers and their reading habits, making sure that each referral is uniquely customized to their needs.

With our individualized publication referrals, you'll never ever have to lose your time or money on a publication that doesn't reverberate with you. Let us help you uncover your next literary journey and enhance your analysis experience.

GIFT THE PLEASURE OF READING APACHE SPARK IN 24 HOURS SAMS TEACH YOURSELF

Publications make remarkable gifts for any event. At our shop, we offer a broad choice of books for any ages and passions, making sure that you'll locate the ideal publication for your enjoyed ones.

Absolutely nothing defeats the feeling of giving someone the **pleasure of reviewing Apache Spark In 24 Hours Sams Teach Yourself**. Whether it's an awesome mystery, a heartwarming

love, or an exciting biography, publications have the power to deliver viewers to brand-new worlds and trigger their creative imaginations.

With our tailored recommendation service, we can help you choose the perfect publication for your recipient. Our educated team will take into account their interests and reviewing choices to supply tailored suggestions that are sure to pleasure.

Unsure which book to select? Take into consideration among our **present** cards, which enables the recipient to discover our collection and choose a publication that speaks with them directly.

So why not share the magic of reading with those you enjoy? Provide the **Apache Spark In 24 Hours Sams Teach Yourself** of a book and make their day a little better!

STAY UPGRADED WITH THE MOST UP TO DATE RELEASES AND BESTSELLERS

At our shop, we pride ourselves on maintaining our collection up-to-date with the **Apache Spark In 24 Hours Sams Teach Yourself latest releases** and **bestsellers**. We understand just how important it is to remain in the loophole with the literary world and have the most recent titles at your fingertips.

Whether you're a follower of secret novels, self-help guides, or love stories, we have the **latest launches** and **bestsellers** in every genre. Our group stays on top of brand-new releases and popular titles, so you can always discover something to read that's fresh and interesting.

Searching for something certain? Our knowledgeable personnel can aid you track down hard-to-find titles and advise similar authors and titles you may appreciate. Whether you're a seasoned visitor or brand-new to the literary world, we have something for everybody.

What's New

A few of our current additions to the collection consist of:

These titles are quickly ending up being brand-new fan favorites and are sure to astound you with their gripping storytelling and unique perspectives.

Do not miss out on the **Apache Spark In 24 Hours Sams Teach Yourself most current launches** and bestsellers - shop with us today and uncover your following favorite read!

TAKE PLEASURE IN SPECIAL DISCOUNTS AND PROMOS

At our shop, our team believe that everyone should have accessibility to excellent books at budget-friendly prices. That's why we provide **Apache Spark In 24 Hours Sams Teach Yourself unique discounts** and **promotions** to our valued customers.

Register for our newsletter to get updates on the most recent **promotions** and discount rates offered. You can likewise follow us on social networks, where we regularly introduce special deals and limited-time bargains.

As a part of our loyalty program, you'll obtain even more **special discounts** and rewards, consisting

of free shipping and special grant purchase.

Whether you're a brand-new consumer or a loyal consumer, we intend to see to it that you have access to the best feasible bargains on our books. Shop with us today and make use of our **exclusive discount rates** and **promotions**!

PRACTICAL ONLINE BUYING AND QUICK SHIPMENT

Purchasing your favored publications has never been less complicated! At our store, we offer a **convenient on-line purchasing** system that allows you to surf our comprehensive collection from the comfort of your home. Just add your preferred books to your cart, enter your delivery details, and proceed to checkout. It's that easy!

But the convenience doesn't quit there. We also offer **quick delivery** alternatives that guarantee you receive your publications in a timely way. Whether you need them for a book club conference or simply can't wait to study a new story, we have actually got you covered.

Shipping Choices

- Criterion shipment: Our conventional distribution option typically takes 3-7 organization days to arrive.
- Expedited distribution: Need your publications quicker? Our expedited delivery option guarantees delivery within 1-3 company days.
- Curbside pickup: If you choose to pick up your books face to face, we likewise provide curbside pick-up at pick places.

Our group works hard to guarantee your Apache Spark In 24 Hours Sams Teach Yourself orders are refined and delivered as swiftly as feasible. And also, with our easy-to-use online tracking system, you can watch on your bundle's development every step of the method.

Experience the benefit and speed of our on-line buying and delivery system. Position your order today and take the very first step on your next literary experience!

CLIENT COMPLETE SATISFACTION ENSURED

At our shop, we take terrific satisfaction in making sure that our consumers are pleased with their purchasing experience. From the moment you begin searching our site to the last shipment of your order, we are committed to offering phenomenal solution.

If you have any type of inquiries or worries regarding your order, please do not hesitate to connect to our customer service group. We are below to aid you and make certain that you are entirely satisfied with your acquisition.

We provide a hassle-free return plan for any type of items that do not meet your expectations. We additionally offer a secure settlement portal, so you can shop with confidence, understanding that your personal information is safeguarded.

We stand behind the top quality of Apache Spark In 24 Hours Sams Teach Yourself publications and are confident that you will be satisfied with your purchase. If you are not completely satisfied with your order, we will do whatever we can to make it right.

Our commitment to consumer contentment is steady, and we aim to surpass your assumptions in every method feasible. Shop with us today and experience the difference!

FINAL THOUGHT

We wish you're as delighted about the globe of Apache Spark In 24 Hours Sams Teach Yourself books as we are! Our shop is devoted to providing a remarkable analysis experience for all. With our wide range of books, from new launches to classic standards, there is something for everyone. Our dedication to client satisfaction implies that you can shop confidently, recognizing that your order remains in good hands.

Do not neglect to connect with us and fellow publication enthusiasts in our vibrant area. Share your preferred reads, take part in thoughtful conversations, and find brand-new literary treasures. And if you ever need assistance finding the best publication, our experienced team is here to give personalized suggestions.

Apache Spark in 24 Hours, Sams Teach Yourself *Sams Publishing*

Apache Spark is a fast, scalable, and flexible open source distributed processing engine for big data systems and is one of the most active open source big data projects to date. In just 24 lessons of one hour or less, Sams Teach Yourself Apache Spark in 24 Hours helps you build practical Big Data solutions that leverage Spark's amazing speed, scalability, simplicity, and versatility. This book's straightforward, step-by-step approach shows you how to deploy, program, optimize, manage, integrate, and extend Spark—now, and for years to come. You'll discover how to create powerful solutions encompassing cloud computing, real-time stream processing, machine learning, and more. Every lesson builds on what you've already learned, giving you a rock-solid foundation for real-world success. Whether you are a data analyst, data engineer, data scientist, or data steward, learning Spark will help you to advance your career or embark on a new career in the booming area of Big Data. Learn how to • Discover what Apache Spark does and how it fits into the Big Data landscape • Deploy and run Spark locally or in the cloud • Interact with Spark from the shell • Make the most of the Spark Cluster Architecture • Develop Spark applications with Scala and functional Python • Program with the Spark API, including transformations and actions • Apply practical data engineering/analysis approaches designed for Spark • Use Resilient Distributed Datasets (RDDs) for caching, persistence, and output • Optimize Spark solution performance • Use Spark with SQL (via Spark SQL) and with NoSQL (via Cassandra) • Leverage cutting-edge functional programming techniques • Extend Spark with streaming, R, and Sparkling Water • Start building Spark-based machine learning and graph-processing applications • Explore advanced messaging technologies, including Kafka • Preview and prepare for Spark's next generation of innovations Instructions walk you through common questions, issues, and tasks; Q-and-As, Quizzes, and Exercises build and test

your knowledge; "Did You Know?" tips offer insider advice and shortcuts; and "Watch Out!" alerts help you avoid pitfalls. By the time you're finished, you'll be comfortable using Apache Spark to solve a wide spectrum of Big Data problems.

Apache Spark in 24 Hours

Learning Spark *O'Reilly Media*

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

Sams Teach Yourself Hadoop in 24 Hours

Data Analytics with Spark Using Python *Addison-Wesley Professional*

Solve Data Analytics Problems with Spark, PySpark, and Related Open Source Tools Spark is at the heart of today's Big Data revolution, helping data professionals supercharge efficiency and performance in a wide range of data processing and analytics tasks. In this guide, Big Data expert Jeffrey Aven covers all you need to know to leverage Spark, together with its extensions, subprojects, and wider ecosystem. Aven combines a language-agnostic introduction to foundational Spark concepts with extensive programming examples utilizing the popular and intuitive PySpark development environment. This guide's focus on Python makes it widely accessible to large audiences of data professionals, analysts, and developers—even those with little Hadoop or Spark experience. Aven's broad coverage ranges from basic to advanced Spark programming, and Spark SQL to machine learning. You'll learn how to efficiently manage all forms of data with Spark: streaming, structured, semi-structured, and unstructured. Throughout, concise topic overviews quickly get you up to speed, and extensive hands-on exercises prepare you to solve real problems. Coverage includes: • Understand Spark's evolving role in the Big Data and Hadoop ecosystems • Create Spark clusters using various deployment modes • Control and optimize the operation of Spark clusters and applications • Master Spark Core RDD API programming techniques • Extend, accelerate, and optimize Spark routines with advanced API platform constructs, including shared variables, RDD storage, and partitioning • Efficiently integrate Spark with both SQL and nonrelational data stores • Perform stream processing and messaging with Spark Streaming and Apache Kafka • Implement predictive modeling with SparkR and Spark MLlib

High Performance Spark

Best Practices for Scaling and Optimizing Apache Spark *"O'Reilly Media, Inc."*

Apache Spark is amazing when everything clicks. But if you haven't seen the performance improvements you expected, or still don't feel confident enough to use Spark in production, this practical book is for you. Authors Holden Karau and Rachel Warren demonstrate performance optimizations to help your Spark queries run faster and handle larger data sizes, while using fewer resources. Ideal for software engineers, data engineers, developers, and system administrators working with large-scale data applications, this book describes techniques that can reduce data infrastructure costs and developer hours. Not only will you gain a more comprehensive understanding of Spark, you'll also learn how to make it sing. With this book, you'll explore: How Spark SQL's new interfaces improve performance over SQL's RDD data structure The choice between data joins in Core Spark and Spark SQL Techniques for getting the most out of standard RDD transformations How to work around performance issues in Spark's key/value pair paradigm Writing high-performance Spark code without Scala or the JVM How to test for functionality and performance when applying suggested improvements Using Spark MLlib and Spark ML machine learning libraries Spark's Streaming components and external community packages

Spark: The Definitive Guide

Big Data Processing Made Simple *"O'Reilly Media, Inc."*

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

Spark in Action

Covers Apache Spark 3 with Examples in Java, Python, and Scala *Simon and Schuster*

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

Hadoop in 24 Hours, Sams Teach Yourself *Sams Publishing*

Apache Hadoop is the technology at the heart of the Big Data revolution, and Hadoop skills are in enormous demand. Now, in just 24 lessons of one hour or less, you can learn all the skills and techniques you'll need to deploy each key component of a Hadoop platform in your local environment or in the cloud, building a fully functional Hadoop cluster and using it with real programs and datasets. Each short, easy lesson builds on all that's come before, helping you master all of Hadoop's essentials, and extend it to meet your unique challenges. Apache Hadoop in 24 Hours, Sams Teach Yourself covers all this, and much more: Understanding Hadoop and the Hadoop Distributed File System (HDFS) Importing data into Hadoop, and process it there Mastering basic MapReduce Java programming, and using advanced MapReduce API concepts Making the most of Apache Pig and Apache Hive Implementing and administering YARN Taking advantage of the full Hadoop ecosystem Managing Hadoop clusters with Apache Ambari Working with the Hadoop User Environment (HUE) Scaling, securing, and troubleshooting Hadoop environments Integrating Hadoop into the enterprise Deploying Hadoop in the cloud Getting started with Apache Spark Step-by-step

instructions walk you through common questions, issues, and tasks; Q-and-As, Quizzes, and Exercises build and test your knowledge; "Did You Know?" tips offer insider advice and shortcuts; and "Watch Out!" alerts help you avoid pitfalls. By the time you're finished, you'll be comfortable using Apache Hadoop to solve a wide spectrum of Big Data problems.

Learning PySpark *Packt Publishing Ltd*

Build data-intensive applications locally and deploy at scale using the combined powers of Python and Spark 2.0 About This Book Learn why and how you can efficiently use Python to process data and build machine learning models in Apache Spark 2.0 Develop and deploy efficient, scalable real-time Spark solutions Take your understanding of using Spark with Python to the next level with this jump start guide Who This Book Is For If you are a Python developer who wants to learn about the Apache Spark 2.0 ecosystem, this book is for you. A firm understanding of Python is expected to get the best out of the book. Familiarity with Spark would be useful, but is not mandatory. What You Will Learn Learn about Apache Spark and the Spark 2.0 architecture Build and interact with Spark DataFrames using Spark SQL Learn how to solve graph and deep learning problems using GraphFrames and TensorFrames respectively Read, transform, and understand data and use it to train machine learning models Build machine learning models with MLlib and ML Learn how to submit your applications programmatically using spark-submit Deploy locally built applications to a cluster In Detail Apache Spark is an open source framework for efficient cluster computing with a strong interface for data parallelism and fault tolerance. This book will show you how to leverage the power of Python and put it to use in the Spark ecosystem. You will start by getting a firm understanding of the Spark 2.0 architecture and how to set up a Python environment for Spark. You will get familiar with the modules available in PySpark. You will learn how to abstract data with RDDs and DataFrames and understand the streaming capabilities of PySpark. Also, you will get a thorough overview of machine learning capabilities of PySpark using ML and MLlib, graph processing using GraphFrames, and polyglot persistence using Blaze. Finally, you will learn how to deploy your applications to the cloud using the spark-submit command. By the end of this book, you will have established a firm understanding of the Spark Python API and how it can be used to build data-intensive applications. Style and approach This book takes a very comprehensive, step-by-step approach so you understand how the Spark ecosystem can be used with Python to develop efficient, scalable solutions. Every chapter is standalone and written in a very easy-to-understand manner, with a focus on both the hows and the whys of each concept.

Learning Spark

Lightning-Fast Big Data Analysis *"O'Reilly Media, Inc."*

Data in all domains is getting bigger. How can you work with it efficiently? Recently updated for Spark 1.3, this book introduces Apache Spark, the open source cluster computing system that makes data analytics fast to write and fast to run. With Spark, you can tackle big datasets quickly through simple APIs in Python, Java, and Scala. This edition includes new information on Spark SQL, Spark Streaming, setup, and Maven coordinates. Written by the developers of Spark, this book will

have data scientists and engineers up and running in no time. 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. Quickly dive into Spark capabilities such as distributed datasets, in-memory caching, and the interactive shell Leverage Spark's powerful built-in libraries, including Spark SQL, Spark Streaming, and MLlib Use one programming paradigm instead of mixing and matching tools like Hive, Hadoop, Mahout, and Storm Learn how to deploy interactive, batch, and streaming applications Connect to data sources including HDFS, Hive, JSON, and S3 Master advanced topics like data partitioning and shared variables

Beginning Apache Spark 2

With Resilient Distributed Datasets, Spark SQL, Structured Streaming and Spark Machine Learning library *Apress*

Develop applications for the big data landscape with Spark and Hadoop. This book also explains the role of Spark in developing scalable machine learning and analytics applications with Cloud technologies. Beginning Apache Spark 2 gives you an introduction to Apache Spark and shows you how to work with it. Along the way, you'll discover resilient distributed datasets (RDDs); use Spark SQL for structured data; and learn stream processing and build real-time applications with Spark Structured Streaming. Furthermore, you'll learn the fundamentals of Spark ML for machine learning and much more. After you read this book, you will have the fundamentals to become proficient in using Apache Spark and know when and how to apply it to your big data applications. What You Will Learn Understand Spark unified data processing platform How to run Spark in Spark Shell or Databricks Use and manipulate RDDs Deal with structured data using Spark SQL through its operations and advanced functions Build real-time applications using Spark Structured Streaming Develop intelligent applications with the Spark Machine Learning library Who This Book Is For Programmers and developers active in big data, Hadoop, and Java but who are new to the Apache Spark platform.

Spark Cookbook *Packt Publishing Ltd*

By introducing in-memory persistent storage, Apache Spark eliminates the need to store intermediate data in filesystems, thereby increasing processing speed by up to 100 times. This book will focus on how to analyze large and complex sets of data. Starting with installing and configuring Apache Spark with various cluster managers, you will cover setting up development environments. You will then cover various recipes to perform interactive queries using Spark SQL and real-time streaming with various sources such as Twitter Stream and Apache Kafka. You will then focus on machine learning, including supervised learning, unsupervised learning, and recommendation engine algorithms. After mastering graph processing using GraphX, you will cover various recipes for cluster optimization and troubleshooting.

Big Data Analytics with Spark

A Practitioner's Guide to Using Spark for Large Scale Data Analysis *Apress*

Big Data Analytics with Spark is a step-by-step guide for learning Spark, which is an open-source fast and general-purpose cluster computing framework for large-scale data analysis. You will learn how to use Spark for different types of big data analytics projects, including batch, interactive, graph, and stream data analysis as well as machine learning. In addition, this book will help you become a much sought-after Spark expert. Spark is one of the hottest Big Data technologies. The amount of data generated today by devices, applications and users is exploding. Therefore, there is a critical need for tools that can analyze large-scale data and unlock value from it. Spark is a powerful technology that meets that need. You can, for example, use Spark to perform low latency computations through the use of efficient caching and iterative algorithms; leverage the features of its shell for easy and interactive Data analysis; employ its fast batch processing and low latency features to process your real time data streams and so on. As a result, adoption of Spark is rapidly growing and is replacing Hadoop MapReduce as the technology of choice for big data analytics. This book provides an introduction to Spark and related big-data technologies. It covers Spark core and its add-on libraries, including Spark SQL, Spark Streaming, GraphX, and MLlib. Big Data Analytics with Spark is therefore written for busy professionals who prefer learning a new technology from a consolidated source instead of spending countless hours on the Internet trying to pick bits and pieces from different sources. The book also provides a chapter on Scala, the hottest functional programming language, and the program that underlies Spark. You'll learn the basics of functional programming in Scala, so that you can write Spark applications in it. What's more, Big Data Analytics with Spark provides an introduction to other big data technologies that are commonly used along with Spark, like Hive, Avro, Kafka and so on. So the book is self-sufficient; all the technologies that you need to know to use Spark are covered. The only thing that you are expected to know is programming in any language. There is a critical shortage of people with big data expertise, so companies are willing to pay top dollar for people with skills in areas like Spark and Scala. So reading this book and absorbing its principles will provide a boost—possibly a big boost—to your career.

Hadoop For Dummies *John Wiley & Sons*

Let Hadoop For Dummies help harness the power of your data and rein in the information overload. Big data has become big business, and companies and organizations of all sizes are struggling to find ways to retrieve valuable information from their massive data sets with becoming overwhelmed. Enter Hadoop and this easy-to-understand For Dummies guide. Hadoop For Dummies helps readers understand the value of big data, make a business case for using Hadoop, navigate the Hadoop ecosystem, and build and manage Hadoop applications and clusters. Explains the origins of Hadoop, its economic benefits, and its functionality and practical applications. Helps you find your way around the Hadoop ecosystem, program MapReduce, utilize design patterns, and get your Hadoop cluster up and running quickly and easily. Details how to use Hadoop applications for data mining, web analytics and personalization, large-scale text processing, data science, and problem-solving. Shows you how to improve the value of your Hadoop cluster, maximize your investment in Hadoop, and avoid common pitfalls when building your Hadoop cluster. From programmers challenged with

building and maintaining affordable, scalable data systems to administrators who must deal with huge volumes of information effectively and efficiently, this how-to has something to help you with Hadoop.

Introduction to Apache Flink

Stream Processing for Real Time and Beyond *O'Reilly Media, Inc.*

There's growing interest in learning how to analyze streaming data in large-scale systems such as web traffic, financial transactions, machine logs, industrial sensors, and many others. But analyzing data streams at scale has been difficult to do well—until now. This practical book delivers a deep introduction to Apache Flink, a highly innovative open source stream processor with a surprising range of capabilities. Authors Ellen Friedman and Kostas Tzoumas show technical and nontechnical readers alike how Flink is engineered to overcome significant tradeoffs that have limited the effectiveness of other approaches to stream processing. You'll also learn how Flink has the ability to handle both stream and batch data processing with one technology. Learn the consequences of not doing streaming well—in retail and marketing, IoT, telecom, and banking and finance. Explore how to design data architecture to gain the best advantage from stream processing. Get an overview of Flink's capabilities and features, along with examples of how companies use Flink, including in production. Take a technical dive into Flink, and learn how it handles time and stateful computation. Examine how Flink processes both streaming (unbounded) and batch (bounded) data without sacrificing performance.

Frank Kane's Taming Big Data with Apache Spark and Python *Packt Publishing Ltd*

Frank Kane's hands-on Spark training course, based on his bestselling Taming Big Data with Apache Spark and Python video, now available in a book. Understand and analyze large data sets using Spark on a single system or on a cluster. About This Book Understand how Spark can be distributed across computing clusters. Develop and run Spark jobs efficiently using Python. A hands-on tutorial by Frank Kane with over 15 real-world examples teaching you Big Data processing with Spark. Who This Book Is For If you are a data scientist or data analyst who wants to learn Big Data processing using Apache Spark and Python, this book is for you. If you have some programming experience in Python, and want to learn how to process large amounts of data using Apache Spark, Frank Kane's Taming Big Data with Apache Spark and Python will also help you. What You Will Learn Find out how you can identify Big Data problems as Spark problems. Install and run Apache Spark on your computer or on a cluster. Analyze large data sets across many CPUs using Spark's Resilient Distributed Datasets. Implement machine learning on Spark using the MLlib library. Process continuous streams of data in real time using the Spark streaming module. Perform complex network analysis using Spark's GraphX library. Use Amazon's Elastic MapReduce service to run your Spark jobs on a cluster. In Detail Frank Kane's Taming Big Data with Apache Spark and Python is your companion to learning Apache Spark in a hands-on manner. Frank will start you off by teaching you how to set up Spark on a single system or on a cluster, and you'll soon move on to analyzing large data sets using Spark RDD, and developing and running effective Spark jobs quickly using Python. Apache Spark has emerged as the

next big thing in the Big Data domain – quickly rising from an ascending technology to an established superstar in just a matter of years. Spark allows you to quickly extract actionable insights from large amounts of data, on a real-time basis, making it an essential tool in many modern businesses. Frank has packed this book with over 15 interactive, fun-filled examples relevant to the real world, and he will empower you to understand the Spark ecosystem and implement production-grade real-time Spark projects with ease. Style and approach Frank Kane's *Taming Big Data with Apache Spark and Python* is a hands-on tutorial with over 15 real-world examples carefully explained by Frank in a step-by-step manner. The examples vary in complexity, and you can move through them at your own pace.

Big Data Processing with Apache Spark

Efficiently tackle large datasets and big data analysis with Spark and Python *Packt Publishing Ltd*

No need to spend hours ploughing through endless data – let Spark, one of the fastest big data processing engines available, do the hard work for you. Key Features Get up and running with Apache Spark and Python Integrate Spark with AWS for real-time analytics Apply processed data streams to machine learning APIs of Apache Spark Book Description Processing big data in real time is challenging due to scalability, information consistency, and fault-tolerance. This book teaches you how to use Spark to make your overall analytical workflow faster and more efficient. You'll explore all core concepts and tools within the Spark ecosystem, such as Spark Streaming, the Spark Streaming API, machine learning extension, and structured streaming. You'll begin by learning data processing fundamentals using Resilient Distributed Datasets (RDDs), SQL, Datasets, and Dataframes APIs. After grasping these fundamentals, you'll move on to using Spark Streaming APIs to consume data in real time from TCP sockets, and integrate Amazon Web Services (AWS) for stream consumption. By the end of this book, you'll not only have understood how to use machine learning extensions and structured streams but you'll also be able to apply Spark in your own upcoming big data projects. What you will learn Write your own Python programs that can interact with Spark Implement data stream consumption using Apache Spark Recognize common operations in Spark to process known data streams Integrate Spark streaming with Amazon Web Services (AWS) Create a collaborative filtering model with the movielens dataset Apply processed data streams to Spark machine learning APIs Who this book is for Data Processing with Apache Spark is for you if you are a software engineer, architect, or IT professional who wants to explore distributed systems and big data analytics. Although you don't need any knowledge of Spark, prior experience of working with Python is recommended.

Spark in Action *Simon and Schuster*

Summary Spark in Action teaches you the theory and skills you need to effectively handle batch and streaming data using Spark. Fully updated for Spark 2.0. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Big data systems distribute datasets across clusters of machines, making it a challenge to efficiently query,

stream, and interpret them. Spark can help. It is a processing system designed specifically for distributed data. It provides easy-to-use interfaces, along with the performance you need for production-quality analytics and machine learning. Spark 2 also adds improved programming APIs, better performance, and countless other upgrades. About the Book Spark in Action teaches you the theory and skills you need to effectively handle batch and streaming data using Spark. You'll get comfortable with the Spark CLI as you work through a few introductory examples. Then, you'll start programming Spark using its core APIs. Along the way, you'll work with structured data using Spark SQL, process near-real-time streaming data, apply machine learning algorithms, and munge graph data using Spark GraphX. For a zero-effort startup, you can download the preconfigured virtual machine ready for you to try the book's code. What's Inside Updated for Spark 2.0 Real-life case studies Spark DevOps with Docker Examples in Scala, and online in Java and Python About the Reader Written for experienced programmers with some background in big data or machine learning. About the Authors Petar Zečević and Marko Bonać are seasoned developers heavily involved in the Spark community. Table of Contents PART 1 - FIRST STEPS Introduction to Apache Spark Spark fundamentals Writing Spark applications The Spark API in depth PART 2 - MEET THE SPARK FAMILY Sparkling queries with Spark SQL Ingesting data with Spark Streaming Getting smart with MLlib ML: classification and clustering Connecting the dots with GraphX PART 3 - SPARK OPS Running Spark Running on a Spark standalone cluster Running on YARN and Mesos PART 4 - BRINGING IT TOGETHER Case study: real-time dashboard Deep learning on Spark with H2O

Beginning Apache Spark Using Azure Databricks

Unleashing Large Cluster Analytics in the Cloud *Apress*

Analyze vast amounts of data in record time using Apache Spark with Databricks in the Cloud. Learn the fundamentals, and more, of running analytics on large clusters in Azure and AWS, using Apache Spark with Databricks on top. Discover how to squeeze the most value out of your data at a mere fraction of what classical analytics solutions cost, while at the same time getting the results you need, incrementally faster. This book explains how the confluence of these pivotal technologies gives you enormous power, and cheaply, when it comes to huge datasets. You will begin by learning how cloud infrastructure makes it possible to scale your code to large amounts of processing units, without having to pay for the machinery in advance. From there you will learn how Apache Spark, an open source framework, can enable all those CPUs for data analytics use. Finally, you will see how services such as Databricks provide the power of Apache Spark, without you having to know anything about configuring hardware or software. By removing the need for expensive experts and hardware, your resources can instead be allocated to actually finding business value in the data. This book guides you through some advanced topics such as analytics in the cloud, data lakes, data ingestion, architecture, machine learning, and tools, including Apache Spark, Apache Hadoop, Apache Hive, Python, and SQL. Valuable exercises help reinforce what you have learned. What You Will Learn Discover the value of big data analytics that leverage the power of the cloud Get started with Databricks using SQL and Python in either Microsoft Azure or AWS Understand the underlying technology, and how the cloud and Apache Spark fit into the bigger picture See how these tools are

used in the real world Run basic analytics, including machine learning, on billions of rows at a fraction of a cost or free Who This Book Is For Data engineers, data scientists, and cloud architects who want or need to run advanced analytics in the cloud. It is assumed that the reader has data experience, but perhaps minimal exposure to Apache Spark and Azure Databricks. The book is also recommended for people who want to get started in the analytics field, as it provides a strong foundation.

Advanced Analytics with Spark

Patterns for Learning from Data at Scale *O'Reilly Media, Inc."*

In this practical book, four Cloudera data scientists present a set of self-contained patterns for performing large-scale data analysis with Spark. The authors bring Spark, statistical methods, and real-world data sets together to teach you how to approach analytics problems by example. You'll start with an introduction to Spark and its ecosystem, and then dive into patterns that apply common techniques—classification, collaborative filtering, and anomaly detection among others—to fields such as genomics, security, and finance. If you have an entry-level understanding of machine learning and statistics, and you program in Java, Python, or Scala, you'll find these patterns useful for working on your own data applications. Patterns include: Recommending music and the Audioscrobbler data set Predicting forest cover with decision trees Anomaly detection in network traffic with K-means clustering Understanding Wikipedia with Latent Semantic Analysis Analyzing co-occurrence networks with GraphX Geospatial and temporal data analysis on the New York City Taxi Trips data Estimating financial risk through Monte Carlo simulation Analyzing genomics data and the BDG project Analyzing neuroimaging data with PySpark and Thunder

Learning Spark SQL *Packt Publishing Ltd*

Design, implement, and deliver successful streaming applications, machine learning pipelines and graph applications using Spark SQL API About This Book Learn about the design and implementation of streaming applications, machine learning pipelines, deep learning, and large-scale graph processing applications using Spark SQL APIs and Scala. Learn data exploration, data munging, and how to process structured and semi-structured data using real-world datasets and gain hands-on exposure to the issues and challenges of working with noisy and "dirty" real-world data. Understand design considerations for scalability and performance in web-scale Spark application architectures. Who This Book Is For If you are a developer, engineer, or an architect and want to learn how to use Apache Spark in a web-scale project, then this is the book for you. It is assumed that you have prior knowledge of SQL querying. A basic programming knowledge with Scala, Java, R, or Python is all you need to get started with this book. What You Will Learn Familiarize yourself with Spark SQL programming, including working with DataFrame/Dataset API and SQL Perform a series of hands-on exercises with different types of data sources, including CSV, JSON, Avro, MySQL, and MongoDB Perform data quality checks, data visualization, and basic statistical analysis tasks Perform data munging tasks on publically available datasets Learn how to use Spark SQL and Apache Kafka to build streaming applications Learn key performance-tuning tips and tricks in Spark SQL applications

Learn key architectural components and patterns in large-scale Spark SQL applications In Detail In the past year, Apache Spark has been increasingly adopted for the development of distributed applications. Spark SQL APIs provide an optimized interface that helps developers build such applications quickly and easily. However, designing web-scale production applications using Spark SQL APIs can be a complex task. Hence, understanding the design and implementation best practices before you start your project will help you avoid these problems. This book gives an insight into the engineering practices used to design and build real-world, Spark-based applications. The book's hands-on examples will give you the required confidence to work on any future projects you encounter in Spark SQL. It starts by familiarizing you with data exploration and data munging tasks using Spark SQL and Scala. Extensive code examples will help you understand the methods used to implement typical use-cases for various types of applications. You will get a walkthrough of the key concepts and terms that are common to streaming, machine learning, and graph applications. You will also learn key performance-tuning details including Cost Based Optimization (Spark 2.2) in Spark SQL applications. Finally, you will move on to learning how such systems are architected and deployed for a successful delivery of your project. Style and approach This book is a hands-on guide to designing, building, and deploying Spark SQL-centric production applications at scale.

Practical Apache Spark

Using the Scala API *Apress*

Work with Apache Spark using Scala to deploy and set up single-node, multi-node, and high-availability clusters. This book discusses various components of Spark such as Spark Core, DataFrames, Datasets and SQL, Spark Streaming, Spark MLib, and R on Spark with the help of practical code snippets for each topic. Practical Apache Spark also covers the integration of Apache Spark with Kafka with examples. You'll follow a learn-to-do-by-yourself approach to learning – learn the concepts, practice the code snippets in Scala, and complete the assignments given to get an overall exposure. On completion, you'll have knowledge of the functional programming aspects of Scala, and hands-on expertise in various Spark components. You'll also become familiar with machine learning algorithms with real-time usage. What You Will Learn Discover the functional programming features of Scala Understand the complete architecture of Spark and its components Integrate Apache Spark with Hive and Kafka Use Spark SQL, DataFrames, and Datasets to process data using traditional SQL queries Work with different machine learning concepts and libraries using Spark's MLlib packages Who This Book Is For Developers and professionals who deal with batch and stream data processing.

Spark

Big Data Cluster Computing in Production *John Wiley & Sons*

Production-targeted Spark guidance with real-world use cases Spark: Big Data Cluster Computing in Production goes beyond general Spark overviews to provide targeted guidance toward using lightning-fast big-data clustering in production. Written by an expert team well-known in the big data

community, this book walks you through the challenges in moving from proof-of-concept or demo Spark applications to live Spark in production. Real use cases provide deep insight into common problems, limitations, challenges, and opportunities, while expert tips and tricks help you get the most out of Spark performance. Coverage includes Spark SQL, Tachyon, Kerberos, ML Lib, YARN, and Mesos, with clear, actionable guidance on resource scheduling, db connectors, streaming, security, and much more. Spark has become the tool of choice for many Big Data problems, with more active contributors than any other Apache Software project. General introductory books abound, but this book is the first to provide deep insight and real-world advice on using Spark in production. Specific guidance, expert tips, and invaluable foresight make this guide an incredibly useful resource for real production settings. Review Spark hardware requirements and estimate cluster size Gain insight from real-world production use cases Tighten security, schedule resources, and fine-tune performance Overcome common problems encountered using Spark in production Spark works with other big data tools including MapReduce and Hadoop, and uses languages you already know like Java, Scala, Python, and R. Lightning speed makes Spark too good to pass up, but understanding limitations and challenges in advance goes a long way toward easing actual production implementation. Spark: Big Data Cluster Computing in Production tells you everything you need to know, with real-world production insight and expert guidance, tips, and tricks.

Big Data Processing with Apache Spark *Lulu.com*

Apache Spark 2.x Cookbook *Packt Publishing Ltd*

Over 70 recipes to help you use Apache Spark as your single big data computing platform and master its libraries About This Book This book contains recipes on how to use Apache Spark as a unified compute engine Cover how to connect various source systems to Apache Spark Covers various parts of machine learning including supervised/unsupervised learning & recommendation engines Who This Book Is For This book is for data engineers, data scientists, and those who want to implement Spark for real-time data processing. Anyone who is using Spark (or is planning to) will benefit from this book. The book assumes you have a basic knowledge of Scala as a programming language. What You Will Learn Install and configure Apache Spark with various cluster managers & on AWS Set up a development environment for Apache Spark including Databricks Cloud notebook Find out how to operate on data in Spark with schemas Get to grips with real-time streaming analytics using Spark Streaming & Structured Streaming Master supervised learning and unsupervised learning using MLlib Build a recommendation engine using MLlib Graph processing using GraphX and GraphFrames libraries Develop a set of common applications or project types, and solutions that solve complex big data problems In Detail While Apache Spark 1.x gained a lot of traction and adoption in the early years, Spark 2.x delivers notable improvements in the areas of API, schema awareness, Performance, Structured Streaming, and simplifying building blocks to build better, faster, smarter, and more accessible big data applications. This book uncovers all these features in the form of structured recipes to analyze and mature large and complex sets of data. Starting with installing and configuring Apache Spark with various cluster managers, you will learn to set up development environments. Further on, you will be introduced to working with RDDs,

DataFrames and Datasets to operate on schema aware data, and real-time streaming with various sources such as Twitter Stream and Apache Kafka. You will also work through recipes on machine learning, including supervised learning, unsupervised learning & recommendation engines in Spark. Last but not least, the final few chapters delve deeper into the concepts of graph processing using GraphX, securing your implementations, cluster optimization, and troubleshooting. Style and approach This book is packed with intuitive recipes supported with line-by-line explanations to help you understand Spark 2.x's real-time processing capabilities and deploy scalable big data solutions. This is a valuable resource for data scientists and those working on large-scale data projects.

Apache Spark for the Enterprise: Setting the Business Free *IBM Redbooks*

Analytics is increasingly an integral part of day-to-day operations at today's leading businesses, and transformation is also occurring through huge growth in mobile and digital channels. Enterprise organizations are attempting to leverage analytics in new ways and transition existing analytics capabilities to respond with more flexibility while making the most efficient use of highly valuable data science skills. The recent growth and adoption of Apache Spark as an analytics framework and platform is very timely and helps meet these challenging demands. The Apache Spark environment on IBM z/OS® and Linux on IBM z Systems™ platforms allows this analytics framework to run on the same enterprise platform as the originating sources of data and transactions that feed it. If most of the data that will be used for Apache Spark analytics, or the most sensitive or quickly changing data is originating on z/OS, then an Apache Spark z/OS based environment will be the optimal choice for performance, security, and governance. This IBM® Redpaper™ publication explores the enterprise analytics market, use of Apache Spark on IBM z Systems™ platforms, integration between Apache Spark and other enterprise data sources, and case studies and examples of what can be achieved with Apache Spark in enterprise environments. It is of interest to data scientists, data engineers, enterprise architects, or anybody looking to better understand how to combine an analytics framework and platform on enterprise systems.

Stream Processing with Apache Spark

Mastering Structured Streaming and Spark Streaming *"O'Reilly Media, Inc."*

Before you can build analytics tools to gain quick insights, you first need to know how to process data in real time. With this practical guide, developers familiar with Apache Spark will learn how to put this in-memory framework to use for streaming data. You'll discover how Spark enables you to write streaming jobs in almost the same way you write batch jobs. Authors Gerard Maas and François Garillot help you explore the theoretical underpinnings of Apache Spark. This comprehensive guide features two sections that compare and contrast the streaming APIs Spark now supports: the original Spark Streaming library and the newer Structured Streaming API. Learn fundamental stream processing concepts and examine different streaming architectures Explore Structured Streaming through practical examples; learn different aspects of stream processing in detail Create and operate streaming jobs and applications with Spark Streaming; integrate Spark Streaming with other Spark APIs Learn advanced Spark Streaming techniques, including

approximation algorithms and machine learning algorithms Compare Apache Spark to other stream processing projects, including Apache Storm, Apache Flink, and Apache Kafka Streams

Business Analytics

A Contemporary Approach *Macmillan International Higher Education*

This innovative new textbook, co-authored by an established academic and a leading practitioner, is the first to bring together issues of cloud computing, business intelligence and big data analytics in order to explore how organisations use cloud technology to analyse data and make decisions. In addition to offering an up-to-date exploration of key issues relating to data privacy and ethics, information governance, and the future of analytics, the text describes the options available in deploying analytic solutions to the cloud and draws on real-world, international examples from companies such as Rolls Royce, Lego, Volkswagen and Samsung. Combining academic and practitioner perspectives that are crucial to the understanding of this growing field, Business Analytics acts as an ideal core text for undergraduate, postgraduate and MBA modules on Big Data, Business and Data Analytics, and Business Intelligence, as well as functioning as a supplementary text for modules in Marketing Analytics. The book is also an invaluable resource for practitioners and will quickly enable the next generation of 'Information Builders' within organisations to understand innovative cloud based-analytic solutions.

PySpark Cookbook

Over 60 recipes for implementing big data processing and analytics using Apache Spark and Python *Packt Publishing Ltd*

Combine the power of Apache Spark and Python to build effective big data applications Key Features Perform effective data processing, machine learning, and analytics using PySpark Overcome challenges in developing and deploying Spark solutions using Python Explore recipes for efficiently combining Python and Apache Spark to process data Book Description Apache Spark is an open source framework for efficient cluster computing with a strong interface for data parallelism and fault tolerance. The PySpark Cookbook presents effective and time-saving recipes for leveraging the power of Python and putting it to use in the Spark ecosystem. You'll start by learning the Apache Spark architecture and how to set up a Python environment for Spark. You'll then get familiar with the modules available in PySpark and start using them effortlessly. In addition to this, you'll discover how to abstract data with RDDs and DataFrames, and understand the streaming capabilities of PySpark. You'll then move on to using ML and MLlib in order to solve any problems related to the machine learning capabilities of PySpark and use GraphFrames to solve graph-processing problems. Finally, you will explore how to deploy your applications to the cloud using the spark-submit command. By the end of this book, you will be able to use the Python API for Apache Spark to solve any problems associated with building data-intensive applications. What you will learn Configure a local instance of PySpark in a virtual environment Install and configure Jupyter in local and multi-node environments Create DataFrames from JSON and a dictionary using pyspark.sql Explore

regression and clustering models available in the ML module Use DataFrames to transform data used for modeling Connect to PubNub and perform aggregations on streams Who this book is for The PySpark Cookbook is for you if you are a Python developer looking for hands-on recipes for using the Apache Spark 2.x ecosystem in the best possible way. A thorough understanding of Python (and some familiarity with Spark) will help you get the best out of the book.

Beginning Apache Spark 3

With DataFrame, Spark SQL, Structured Streaming, and Spark Machine Learning Library *Apress*

Take a journey toward discovering, learning, and using Apache Spark 3.0. In this book, you will gain expertise on the powerful and efficient distributed data processing engine inside of Apache Spark; its user-friendly, comprehensive, and flexible programming model for processing data in batch and streaming; and the scalable machine learning algorithms and practical utilities to build machine learning applications. Beginning Apache Spark 3 begins by explaining different ways of interacting with Apache Spark, such as Spark Concepts and Architecture, and Spark Unified Stack. Next, it offers an overview of Spark SQL before moving on to its advanced features. It covers tips and techniques for dealing with performance issues, followed by an overview of the structured streaming processing engine. It concludes with a demonstration of how to develop machine learning applications using Spark MLlib and how to manage the machine learning development lifecycle. This book is packed with practical examples and code snippets to help you master concepts and features immediately after they are covered in each section. After reading this book, you will have the knowledge required to build your own big data pipelines, applications, and machine learning applications. What You Will Learn Master the Spark unified data analytics engine and its various components Work in tandem to provide a scalable, fault tolerant and performant data processing engine Leverage the user-friendly and flexible programming model to perform simple to complex data analytics using dataframe and Spark SQL Develop machine learning applications using Spark MLlib Manage the machine learning development lifecycle using MLflow Who This Book Is For Data scientists, data engineers and software developers.

Build Your Own IoT Platform

Develop a Fully Flexible and Scalable Internet of Things Platform in 24 Hours *Apress*

Discover how every solution in some way related to the IoT needs a platform and how to create that platform. This book is about being agile and reducing time to market without breaking the bank. It is about designing something that you can scale incrementally without having to do a lot of rework and potentially disrupting your current state of the work. So the key questions are: what does it take, how long does it take, and how much does it take to build your own IoT platform? Build Your Own IoT Platform answers these questions and provides you with step-by-step guidance on how to build your own IoT platform. The author bursts the bubble of IoT platforms and highlights what the core of an IoT platform looks like. There are must-haves and there are nice-to-haves; this book will distinguish

the two and focus on how to build the must-haves. Building your own IoT platform is not only the biggest cost saver, but also can be a satisfying learning experience, giving you control over your project. What You Will Learn Architect an interconnected system Develop a flexible architecture Create a redundant communication platform Prioritize system requirements with a bottom-up approach Who This Book Is For IoT developers and development teams in small- to medium-sized companies. Basic to intermediate programming skills are required.

Data Pipelines with Apache Airflow *Simon and Schuster*

"An Airflow bible. Useful for all kinds of users, from novice to expert." - Rambabu Posa, Sai Aashika Consultancy Data Pipelines with Apache Airflow teaches you how to build and maintain effective data pipelines. A successful pipeline moves data efficiently, minimizing pauses and blockages between tasks, keeping every process along the way operational. Apache Airflow provides a single customizable environment for building and managing data pipelines, eliminating the need for a hodgepodge collection of tools, snowflake code, and homegrown processes. Using real-world scenarios and examples, Data Pipelines with Apache Airflow teaches you how to simplify and automate data pipelines, reduce operational overhead, and smoothly integrate all the technologies in your stack. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Data pipelines manage the flow of data from initial collection through consolidation, cleaning, analysis, visualization, and more. Apache Airflow provides a single platform you can use to design, implement, monitor, and maintain your pipelines. Its easy-to-use UI, plug-and-play options, and flexible Python scripting make Airflow perfect for any data management task. About the book Data Pipelines with Apache Airflow teaches you how to build and maintain effective data pipelines. You'll explore the most common usage patterns, including aggregating multiple data sources, connecting to and from data lakes, and cloud deployment. Part reference and part tutorial, this practical guide covers every aspect of the directed acyclic graphs (DAGs) that power Airflow, and how to customize them for your pipeline's needs. What's inside Build, test, and deploy Airflow pipelines as DAGs Automate moving and transforming data Analyze historical datasets using backfilling Develop custom components Set up Airflow in production environments About the reader For DevOps, data engineers, machine learning engineers, and sysadmins with intermediate Python skills. About the author Bas Harens and Julian de Ruyter are data engineers with extensive experience using Airflow to develop pipelines for major companies. Bas is also an Airflow committer. Table of Contents PART 1 - GETTING STARTED 1 Meet Apache Airflow 2 Anatomy of an Airflow DAG 3 Scheduling in Airflow 4 Templating tasks using the Airflow context 5 Defining dependencies between tasks PART 2 - BEYOND THE BASICS 6 Triggering workflows 7 Communicating with external systems 8 Building custom components 9 Testing 10 Running tasks in containers PART 3 - AIRFLOW IN PRACTICE 11 Best practices 12 Operating Airflow in production 13 Securing Airflow 14 Project: Finding the fastest way to get around NYC PART 4 - IN THE CLOUDS 15 Airflow in the clouds 16 Airflow on AWS 17 Airflow on Azure 18 Airflow in GCP

Learn Azure in a Month of Lunches, Second Edition *Manning Publications*

Learn Azure in a Month of Lunches, Second Edition, is a tutorial on writing, deploying, and running

applications in Azure. In it, you'll work through 21 short lessons that give you real-world experience. Each lesson includes a hands-on lab so you can try out and lock in your new skills. Summary You can be incredibly productive with Azure without mastering every feature, function, and service. Learn Azure in a Month of Lunches, Second Edition gets you up and running quickly, teaching you the most important concepts and tasks in 21 practical bite-sized lessons. As you explore the examples, exercises, and labs, you'll pick up valuable skills immediately and take your first steps to Azure mastery! This fully revised new edition covers core changes to the Azure UI, new Azure features, Azure containers, and the upgraded Azure Kubernetes Service. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Microsoft Azure is vast and powerful, offering virtual servers, application templates, and prebuilt services for everything from data storage to AI. To navigate it all, you need a trustworthy guide. In this book, Microsoft engineer and Azure trainer Iain Foulds focuses on core skills for creating cloud-based applications. About the book Learn Azure in a Month of Lunches, Second Edition, is a tutorial on writing, deploying, and running applications in Azure. In it, you'll work through 21 short lessons that give you real-world experience. Each lesson includes a hands-on lab so you can try out and lock in your new skills. What's inside Understanding Azure beyond point-and-click Securing applications and data Automating your environment Azure services for machine learning, containers, and more About the reader This book is for readers who can write and deploy simple web or client/server applications. About the author Iain Foulds is an engineer and senior content developer with Microsoft. Table of Contents PART 1 - AZURE CORE SERVICES 1 Before you begin 2 Creating a virtual machine 3 Azure Web Apps 4 Introduction to Azure Storage 5 Azure Networking basics PART 2 - HIGH AVAILABILITY AND SCALE 6 Azure Resource Manager 7 High availability and redundancy 8 Load-balancing applications 9 Applications that scale 10 Global databases with Cosmos DB 11 Managing network traffic and routing 12 Monitoring and troubleshooting PART 3 - SECURE BY DEFAULT 13 Backup, recovery, and replication 14 Data encryption 15 Securing information with Azure Key Vault 16 Azure Security Center and updates PART 4 - THE COOL STUFF 17 Machine learning and artificial intelligence 18 Azure Automation 19 Azure containers 20 Azure and the Internet of Things 21 Serverless computing

Big Data Analytics with Java *Packt Publishing Ltd*

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 Data Engine for Hadoop and Spark *IBM Redbooks*

This IBM® Redbooks® publication provides topics to help the technical community take advantage of the resilience, scalability, and performance of the IBM Power Systems™ platform to implement or integrate an IBM Data Engine for Hadoop and Spark solution for analytics solutions to access, manage, and analyze data sets to improve business outcomes. This book documents topics to demonstrate and take advantage of the analytics strengths of the IBM POWER8® platform, the IBM analytics software portfolio, and selected third-party tools to help solve customer's data analytic workload requirements. This book describes how to plan, prepare, install, integrate, manage, and show how to use the IBM Data Engine for Hadoop and Spark solution to run analytic workloads on IBM POWER8. In addition, this publication delivers documentation to complement available IBM analytics solutions to help your data analytic needs. This publication strengthens the position of IBM analytics and big data solutions with a well-defined and documented deployment model within an IBM POWER8 virtualized environment so that customers have a planned foundation for security, scaling, capacity, resilience, and optimization for analytics workloads. This book is targeted at technical professionals (analytics consultants, technical support staff, IT Architects, and IT Specialists) that are responsible for delivering analytics solutions and support on IBM Power Systems.

Data Science on AWS *"O'Reilly Media, Inc."*

With this practical book, AI and machine learning practitioners will learn how to successfully build and deploy data science projects on Amazon Web Services. The Amazon AI and machine learning stack unifies data science, data engineering, and application development to help level up your skills. This guide shows you how to build and run pipelines in the cloud, then integrate the results into applications in minutes instead of days. Throughout the book, authors Chris Fregly and Antje Barth

demonstrate how to reduce cost and improve performance. Apply the Amazon AI and ML stack to real-world use cases for natural language processing, computer vision, fraud detection, conversational devices, and more Use automated machine learning to implement a specific subset of use cases with SageMaker Autopilot Dive deep into the complete model development lifecycle for a BERT-based NLP use case including data ingestion, analysis, model training, and deployment Tie everything together into a repeatable machine learning operations pipeline Explore real-time ML, anomaly detection, and streaming analytics on data streams with Amazon Kinesis and Managed Streaming for Apache Kafka Learn security best practices for data science projects and workflows including identity and access management, authentication, authorization, and more

Data Algorithms

Recipes for Scaling Up with Hadoop and Spark *"O'Reilly Media, Inc."*

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

Machine Learning with Spark *Packt Pub Limited*

Pro Spark Streaming

The Zen of Real-Time Analytics Using Apache Spark *Apress*

Learn the right cutting-edge skills and knowledge to leverage Spark Streaming to implement a wide array of real-time, streaming applications. This book walks you through end-to-end real-time application development using real-world applications, data, and code. Taking an application-first approach, each chapter introduces use cases from a specific industry and uses publicly available datasets from that domain to unravel the intricacies of production-grade design and implementation. The domains covered in Pro Spark Streaming include social media, the sharing economy, finance, online advertising, telecommunication, and IoT. In the last few years, Spark has become synonymous with big data processing. DStreams enhance the underlying Spark processing engine to support streaming analysis with a novel micro-batch processing model. Pro Spark

Streaming by Zubair Nabi will enable you to become a specialist of latency sensitive applications by leveraging the key features of DStreams, micro-batch processing, and functional programming. To this end, the book includes ready-to-deploy examples and actual code. Pro Spark Streaming will act as the bible of Spark Streaming. What You'll Learn Discover Spark Streaming application development and best practices Work with the low-level details of discretized streams Optimize production-grade deployments of Spark Streaming via configuration recipes and instrumentation using Graphite, collectd, and Nagios Ingest data from disparate sources including MQTT, Flume,

Kafka, Twitter, and a custom HTTP receiver Integrate and couple with HBase, Cassandra, and Redis Take advantage of design patterns for side-effects and maintaining state across the Spark Streaming micro-batch model Implement real-time and scalable ETL using data frames, SparkSQL, Hive, and SparkR Use streaming machine learning, predictive analytics, and recommendations Mesh batch processing with stream processing via the Lambda architecture Who This Book Is For Data scientists, big data experts, BI analysts, and data architects.