Apache Hadoop Yarn Moving Beyond Mapreduce And Batch Processing With Apache Hadoop 2 Addison Wesley Data Analytics

Many corporations are finding that the size of their data sets are outgrowing the capability of their systems to store and process them. The data is becoming too big to manage and use with traditional tools. The solution: implementing a big data system. As Big Data Made Easy: A Working Guide to the Complete Hadoop Toolset shows, Apache Hadoop offers a scalable, fault-tolerant system for storing and processing data in parallel. It has a very rich toolset that allows for storage (Hadoop), configuration (YARN and ZooKeeper), collection (Nutch and Solr), processing (Storm, Pig, and Map Reduce), scheduling (Oozie), moving (Sqoop and Avro), monitoring (Chukwa, Ambari, and Hue), testing (Big Top), and analysis (Hive). The problem is that the Internet offers IT pros wading into big data many versions of the truth and some outright falsehoods born of ignorance. What is needed is a book just like this one: a wide-ranging but easily understood set of instructions to explain where to get Hadoop tools, what they can do, how to install them, how to configure them, how to integrate them, and how to use them successfully. And you need an expert who has worked in this area for a decade—someone just like author and big data expert Mike Frampton. Big Data Made Easy approaches the problem of managing massive data sets from a systems perspective, and it explains the roles for each project (like architect and tester, for example) and shows how the Hadoop toolset can be used at each system stage. It explains, in an easily understood manner and through numerous examples, how to use each tool. The book also explains the sliding scale of tools available depending upon data size and when and how to use them. Big Data Made Easy shows developers and architects, as well as testers and project managers, how to: Store big data Configure big data Process big data Schedule processes Move data among SQL and NoSQL systems Monitor data Perform big data analytics Report on big data processes and projects Test big data systems Big Data Made Easy also explains the best part, which is that this toolset is free. Anyone can download it and—with the help of this book—start to use it within a day. With the skills this book will teach you under your belt, you will add value to your company or client immediately, not to mention your career.

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.

Until recently, Hadoop deployments existed on hardware owned and run by organizations. Now, of course, you can acquire the computing resources and network connectivity to run Hadoop clusters in the cloud. But there’s a lot more to deploying

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

This book constitutes the thoroughly refereed post-conference proceedings of the 23rd International Workshop on Job Scheduling Strategies for Parallel Processing, JSSPP 2020, held in New Orleans, LA, USA, in May 2020.* The 6 revised full papers presented were carefully reviewed and selected from 8 submissions. In addition to this, one invited paper and one keynote pare were included in the workshop. The papers cover topics within the fields of resource management and scheduling. They focus on several interesting problems such as resource contention and workload interference, new scheduling policy, scheduling ultrasound simulation workflows, and walltime prediction. * The conference was held virtually due to the COVID-19 pandemic.

Due to the increasing availability of affordable internet services, the number of users, and the need for a wider range of multimedia-based applications, internet usage is on the rise. With so many users and such a large amount of data, the requirements of analyzing large data sets leads to the need for further advancements to information processing. Big Data Processing With Hadoop is an essential reference source that discusses possible solutions for millions of users working with a variety of data applications, who expect fast turnaround responses, but encounter issues with processing data at the rate it comes in. Featuring research on topics such as market basket analytics, scheduler load simulator, and writing YARN applications, this book is ideally designed for IoT professionals, students, and engineers seeking coverage on many of the real-world challenges regarding big data.

The two-volume set LNCS 9952 and LNCS 9953 constitutes the refereed proceedings of the 7th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation, ISoLA 2016, held in Imperial, Corfu, Greece, in October 2016. The papers presented in this volume were carefully reviewed and selected for inclusion in the proceedings. Featuring a track introduction to each section, the papers are organized in topical sections named: statistical model checking; evaluation and reproducibility of program analysis and verification; ModSyn-PP: modular synthesis of programs and processes; semantic heterogeneity in the formal development of complex systems; static and runtime verification: competitors or friends?; rigorous engineering of collective adaptive systems; correctness-by-construction and post-hoc verification: friends or foes?; privacy and security issues in information systems; towards a unified view of modeling and programming; formal methods and safety certification: challenges in the railways domain; RVE: runtime verification and enforcement, the (industrial) application perspective; variability modeling for scalable software evolution; detecting and understanding software doping; learning systems: machine-learning in software products and learning-based analysis of software systems; testing the internet of things; doctoral symposium; industrial track; RERS challenge; and STRESS.

This two volume set (CCIS 623 and 634) constitutes the refereed proceedings of the Second International Conference of Young Computer Scientists, Engineers and Educators, ICYCSEE 2016, held in

Harbin, China, in August 2016. The 91 revised full papers presented were carefully reviewed and selected from 338 submissions. The papers are organized in topical sections on Research Track (Part I) and Education Track, Industry Track, and Demo Track (Part II) and cover a wide range of topics related to social computing, social media, social network analysis, social modeling, social recommendation, machine learning, data mining. This book examines the Internet of Things (IoT) and Data Analytics from a technical, application, and business point of view. Internet of Things and Data Analytics Handbook describes essential technical knowledge, building blocks, processes, design principles, implementation, and marketing for IoT projects. It provides readers with knowledge in planning, designing, and implementing IoT projects. The book is written by experts on the subject matter, including international experts from nine countries in the consumer and enterprise fields of IoT. The text starts with an overview and anatomy of IoT, ecosystem of IoT, communication protocols, networking, and available hardware, both present and future applications and transformations, and business models. The text also addresses big data analytics, machine learning, cloud computing, and consideration of sustainability that are essential to be both socially responsible and successful. Design and implementation processes are illustrated with best practices and case studies in action. In addition, the book: Examines cloud computing, data analytics, and sustainability and how they relate to IoT overs the scope of consumer, government, and enterprise applications Includes best practices, business model, and real-world case studies Hwaiyu Geng, P.E., is a consultant with Amica Research (www.AmicaResearch.org, Palo Alto, California), promoting green planning, design, and construction projects. He has had over 40 years of manufacturing and management experience, working with Westinghouse, Applied Materials, Hewlett Packard, and Intel on multi-million high-tech projects. He has written and presented numerous technical papers at international conferences. Mr. Geng, a patent holder, is also the editor/author of Data Center Handbook (Wiley, 2015).

???? ??? ???? ???? ?? ??? ?? ? ?? ??? ???? ??? ??? ??? ??? ????? ?? ??? ???, ??? ????, ???? ?? ???? ???? ??. ??? ????? ? ??? ????, ?? ??? ????? ???, ??? ?? ??? ????. ??? ??? ??? ? ??? ????? ?? ?? ????? ??? ????.

Archiving has become an increasingly complex process. The challenge is no longer how to store the data but how to store it intelligently, in order to exploit it over time, while maintaining its integrity and authenticity. Digital technologies bring about major transformations, not only in terms of the types of documents that are transferred to and stored in archives, in the behaviors and practices of the humanities and social sciences (digital humanities), but also in terms of the volume of data and the technological capacity for managing and preserving archives (Big Data). Archives in The Digital Age focuses on the impact of these various digital transformations on archives, and examines how the right to memory and the information of future generations is confronted with the right to be forgotten; a digital prerogative that guarantees individuals their private lives and freedoms.

Unlock the power of your data with Hadoop 2.X ecosystem and its data warehousing techniques across large data sets About This Book Conquer the mountain of data using Hadoop 2.X tools The authors succeed in creating a context for Hadoop and its ecosystem Hands-on examples and recipes giving the bigger picture and helping you to master Hadoop 2.X data processing platforms Overcome the challenging data processing problems using this exhaustive course with Hadoop 2.X Who This Book Is For This course is for Java developers, who know scripting, wanting a career shift to Hadoop - Big Data segment of the IT industry. So if you are a novice in Hadoop or an expert, this book will make you reach the most advanced level in Hadoop 2.X. What You Will Learn Best practices for setup and configuration of Hadoop clusters, tailoring the system to the problem at hand Integration with relational databases, using Hive for SQL queries and Sqoop for data transfer Installing and maintaining Hadoop 2.X cluster and its ecosystem Advanced Data Analysis using the Hive, Pig, and Map Reduce programs Machine learning principles with libraries such as Mahout and Batch and Stream data processing using Apache Spark Understand the changes involved in the process in the move from Hadoop 1.0 to

Hadoop 2.0 Dive into YARN and Storm and use YARN to integrate Storm with Hadoop Deploy Hadoop on Amazon Elastic MapReduce and Discover HDFS replacements and learn about HDFS Federation In Detail As Marc Andreessen has said “Data is eating the world,” which can be witnessed today being the age of Big Data, businesses are producing data in huge volumes every day and this rise in tide of data need to be organized and analyzed in a more secured way. With proper and effective use of Hadoop, you can build new-improved models, and based on that you will be able to make the right decisions. The first module, Hadoop beginners Guide will walk you through on understanding Hadoop with very detailed instructions and how to go about using it. Commands are explained using sections called “What just happened” for more clarity and understanding. The second module, Hadoop Real World Solutions Cookbook, 2nd edition, is an essential tutorial to effectively implement a big data warehouse in your business, where you get detailed practices on the latest technologies such as YARN and Spark. Big data has become a key basis of competition and the new waves of productivity growth. Hence, once you get familiar with the basics and implement the end-to-end big data use cases, you will start exploring the third module, Mastering Hadoop. So, now the question is if you need to broaden your Hadoop skill set to the next level after you nail the basics and the advance concepts, then this course is indispensable. When you finish this course, you will be able to tackle the real-world scenarios and become a big data expert using the tools and the knowledge based on the various step-by-step tutorials and recipes. Style and approach This course has covered everything right from the basic concepts of Hadoop till you master the advance mechanisms to become a big data expert. The goal here is to help you learn the basic essentials using the step-by-step tutorials and from there moving toward the recipes with various real-world solutions for you. It covers all the important aspects of Hadoop from system designing and configuring Hadoop, machine learning principles with various libraries with chapters illustrated with code fragments and schematic diagrams. This is a compendious course to explore Hadoop from the basics to the most advanced techniques available in Hadoop 2.X.

NOTE: This title is also available as a free eBook on the Microsoft Download Center. It is offered for sale in print format as a convenience. Get a head start evaluating SQL Server 2014 - guided by two experts who have worked with the technology from the earliest beta. Based on Community Technology Preview 2 (CTP2) software, this guide introduces new features and capabilities, with practical insights on how SQL Server 2014 can meet the needs of your business. Get the early, high-level overview you need to begin preparing your deployment now. Coverage includes: SQL Server 2014 Editions and engine enhancements Mission-critical performance enhancements Hybrid cloud enhancements Self-service Business Intelligence enhancements in Microsoft Excel Enterprise information management enhancements Big Data solutions

Dig deep into the data with a hands-on guide to machine learning Machine Learning: Hands-On for Developers and Technical Professionals provides hands-on instruction and fully-coded working examples for the most common machine learning techniques used by developers and technical professionals. The book contains a breakdown of each ML variant, explaining how it works and how it is used within certain industries, allowing readers to incorporate the presented techniques into their own work as they follow along. A core tenant of machine learning is a strong focus on data preparation, and a full exploration of the various types of learning algorithms illustrates how the proper tools can help any developer extract information and insights from existing data. The book includes a full complement of Instructor's Materials to facilitate use in the classroom, making this resource useful for students and as a professional reference. At its core, machine learning is a mathematical, algorithm-based technology that forms the basis of historical data mining and modern big data science. Scientific analysis of big data requires a working knowledge of machine learning, which forms predictions based on known properties learned from training data. Machine Learning is an accessible, comprehensive guide for the non-mathematician, providing clear guidance that allows readers to: Learn the languages of machine learning including Hadoop, Mahout, and Weka Understand decision trees, Bayesian networks, and artificial neural networks Implement Association Rule, Real Time, and Batch learning Develop a strategic plan for safe, effective, and efficient machine learning By learning to construct a system that can learn from data, readers can increase their utility across industries. Machine learning sits at the core of deep dive data analysis and visualization, which is increasingly in demand as companies discover the goldmine hiding in their existing data. For the tech professional involved in data science, Machine Learning: Hands-On for Developers and Technical Professionals provides the skills and techniques required to dig deeper.

Take a deep dive into the concepts of machine learning as they apply to contemporary business and management. You will learn how machine learning techniques are used to solve fundamental and complex problems in society and industry. Machine Learning for Decision Makers serves as an excellent resource for establishing the relationship of machine learning with IoT, big data, and cognitive and cloud computing to give you an overview of how these modern areas of computing relate to each other. This book introduces a collection of the most important concepts of machine learning and sets them in context with other vital technologies that decision makers need to know about. These concepts span the process from envisioning the problem to applying machine-learning techniques to your particular situation. This discussion also provides an insight to help deploy the results to improve decision-making. The book uses case studies and jargon busting to help you grasp the theory of machine learning quickly. You'll soon gain the big picture of machine learning and how it fits with other cutting-edge IT services. This knowledge will give you confidence in your decisions for the future of your business. What You Will Learn Discover the machine learning, big data, and cloud and cognitive computing technology stack Gain insights into machine learning concepts and practices Understand business and enterprise decision-making using machine learning Absorb machine-learning best practices Who This Book Is For Managers tasked with making key decisions who want to learn how and when machine learning and related technologies can help them.

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

Big Data has been much in the news in recent years, and the advantages conferred by the collection and analysis of large datasets in fields such as marketing, medicine and finance have led to claims that almost any real world problem could be solved if sufficient data were available. This is of course a very simplistic view, and the usefulness of collecting, processing and storing large datasets must always be seen in terms of the communication, processing and storage capabilities of the computing platforms available. This book presents papers from the International Research Workshop, Advanced High Performance Computing Systems, held in Cetraro, Italy, in July 2014. The papers selected for publication here discuss fundamental aspects of the definition of Big Data, as well as considerations from practice where complex datasets are collected, processed and stored. The concepts, problems, methodologies and solutions presented are of much more general applicability than may be suggested by the particular application areas considered. As a result the book will be of interest to all those whose work involves the processing of very large data sets, exascale computing and the emerging fields of data science

The Analytics and Big Data collection offers a “greatest hits” digital compilation of ideas from world-renowned thought leader Thomas Davenport, who helped popularize the terms analytics and big data in the workplace. An agile and prolific thinker, Davenport has written or coauthored more than a dozen bestselling books. Several of these titles are offered together for the first time in this curated digital bundle, including: Big Data at Work, Competing on Analytics, Analytics at Work, and Keeping Up with the Quants. The collection also includes Davenport's popular Harvard Business Review articles, "Data Scientist: The Sexiest Job of the 21st Century" (2012) and "Analytics 3.0" (2013). Combined, these works cover all the bases on analytics and big data: what each term means; the ramifications of each from a technical, consumer, and management perspective; and where each can have the biggest impact on your business. Whether you're an executive, a manager, or a student wanting to learn more, Analytics and Big Data is the most comprehensive collection you'll find on the ever-growing phenomenon of digital data and analysis—and how you can make this rising business trend work for you. Named one of the ten "Masters of the New Economy" by CIO magazine, Thomas Davenport has helped hundreds of companies revitalize their management practices. He combines his interests in research, teaching, and business management as the President's Distinguished Professor of Information Technology & Management at Babson College. Davenport has also taught at Harvard Business School, the University of Chicago, Dartmouth's Tuck School of Business, and the University of Texas at Austin and has directed research centers at Accenture, McKinsey & Company, Ernst & Young, and CSC. He is also an independent Senior Advisor to Deloitte Analytics.

Disorder-assistive and neurotechnological devices are experiencing a boom in the global market. Mounting evidence suggests that approaches based on several different domains should move towards the goal of early diagnosis of individuals affected by neurodevelopmental disorders. Using an interdisciplinary and collaborative approach in diagnosis and support can resolve many hurdles such as lack of awareness, transport, and financial burdens by being made available to individuals at the onset of symptoms. Interdisciplinary Approaches to Altering Neurodevelopmental Disorders is a pivotal reference source that explores neurodevelopmental disorders and a diverse array of diagnostic tools and therapies assisted by neurotechnological devices. While covering a wide range of topics including individual-centered design, artificial intelligence, and multifaceted therapies, this book is ideally designed for neuroscientists, medical practitioners, clinical psychologists, special educators, counselors, therapists, researchers, academicians, and students.

Apache Hadoop YARNMoving beyond MapReduce and Batch Processing with Apache Hadoop 2Addison-Wesley Professional

Get started fast with Apache Hadoop 2, with the first easy, accessible guide to this revolutionary Big Data technology. Building on his unsurpassed experience teaching Hadoop and Big Data, Dr. Douglas Eadline covers all the basics you need to know to install and use Hadoop 2 on both personal computers and servers, and navigate the entire Apache Hadoop ecosystem. Eadline demystifies Hadoop 2, explains the problems it solves, shows how it relates to Big Data, and demonstrates both administrators and users work with it. He explains the central role of MapReduce in Hadoop 1, and how (and why) YARN and Hadoop 2 move beyond MapReduce. You'll find essential information on: Planning and performing Hadoop 2 installations -- including decisions about hardware, software, clustering, and HDFS Using the Hadoop Distributed File System (HDFS) and working around its tradeoffs Running and benchmarking Hadoop 2 programs Working with MapReduce -- including basic programming examples Using higher-level tools, including Pig and Hive Getting started with Apache Hadoop YARN frameworks Administering Hadoop 2 with Ambari, rmadmin, and automated scripts From its Getting Started checklist/flowchart to its roadmap of additional resources, Hadoop 2 Quick-Start Guide is your perfect Hadoop 2 starting point -- and your fastest way to start mastering Big Data.

This book constitutes the proceedings of the 19th International Conference on Fundamental Approaches to Software Engineering, FASE 2016, which took place in Eindhoven, The Netherlands, in April 2016, held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016. The 23 full papers presented in this volume were carefully reviewed and selected from 90 submissions. They were organized in topical sections named: concurrent and distributed systems; model-driven development; analysis and bug triaging; probabilistic and stochastic systems; proof and theorem proving; and verification.

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.

The two-volume set LNICST 236-237 constitutes the post-conference proceedings of the 12th EAI International Conference on Communications and Networking, ChinaCom 2017, held in Xi’an, China, in September 2017. The total of 112 contributions presented in these volumes are carefully reviewed and selected from 178 submissions. Aside from the technical paper sessions the book is organized in topical sections on wireless communications and networking, satellite and space communications and networking, big data network track, multimedia communications and smart networking, signal processing and communications, network and information security, advances and trends of V2X networks.

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 SystemsTM 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® RedpaperTM publication explores the enterprise analytics market, use of Apache Spark on IBM z SystemsTM 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.

Pro Apache Hadoop, Second Edition brings you up to speed on Hadoop – the framework of big data. Revised to cover Hadoop 2.0, the book covers the very latest developments such as YARN (aka MapReduce 2.0), new HDFS high-availability features, and increased scalability in the form of HDFS Federations. All the old content has been revised too, giving the latest on the ins and outs of MapReduce, cluster design, the Hadoop Distributed File System, and more. This book covers everything you need to build your first Hadoop cluster and begin analyzing and deriving value from your business and scientific data. Learn to solve big-data problems the MapReduce way, by breaking a big problem into chunks and creating small-scale solutions that can be flung across thousands upon thousands of nodes to analyze large data volumes in a short amount of wall-clock time. Learn how to let Hadoop take care of distributing and parallelizing your software—you just focus on the code; Hadoop takes care of the rest. Covers all that is new in Hadoop 2.0 Written by a professional involved in Hadoop since day one Takes you quickly to the seasoned pro level on the hottest cloud-computing framework

"Apache Hadoop YARN Fundamentals LiveLessons is the first complete video training course on the basics of Apache Hadoop version 2 with YARN. The tutorial begins with MapReduce and Big Data fundamentals and moves to YARN design, installation (laptop, cluster, and cloud) , administration, running applications (MapReduce2, Pig and Hive), writing new applications, and useful frameworks. Additional coverage of Ambari, Ganglia, Nagios and the Hortonworks HDP is provided."--Resource description page.

This edited volume gathers the proceedings of the Symposium GIS Ostrava 2016, the Rise of Big Spatial Data, held at the Technical University of Ostrava, Czech Republic, March 16–18, 2016. Combining theoretical papers and applications by authors from around the globe, it summarises the latest research findings in the area of big spatial data and key problems related to its utilisation. Welcome to dawn of the big data era: though it’s in sight, it isn’t quite here yet. Big spatial data is characterised by three main features: volume beyond the limit of usual geo-processing, velocity higher than that available using conventional processes, and variety, combining more diverse geodata sources than usual. The popular term denotes a situation in which one or more of these key properties reaches a point at which traditional methods for geodata collection, storage, processing, control, analysis, modelling, validation and visualisation fail to provide effective solutions. >Entering the era of big spatial data calls for finding solutions that address all “small data” issues that soon create “big data” troubles. Resilience for big spatial data means solving the heterogeneity of spatial data sources (in topics, purpose, completeness, guarantee, licensing, coverage etc.), large volumes (from gigabytes to terabytes and more), undue complexity of geo-applications and systems (i.e. combination of standalone applications with web services, mobile platforms and sensor networks), neglected automation of geodata preparation (i.e. harmonisation, fusion), insufficient control of geodata collection and distribution processes (i.e. scarcity and poor quality of metadata and metadata systems), limited analytical tool capacity (i.e. domination of traditional causal-driven analysis), low visual system performance, inefficient knowledge-discovery techniques (for transformation of vast amounts of information into tiny and essential outputs) and much more. These trends are accelerating as sensors become more ubiquitous around the world.

Design, build, and justify an optimal Microsoft IoT footprint to meet your project needs. This book describes common Internet of Things components and architecture and then focuses on Microsoft's Azure components relevant in deploying these solutions. Microsoft-specific topics addressed include: deploying edge devices and pushing intelligence to the edge; connecting IoT devices to Azure and landing data there, applying Azure Machine Learning, analytics, and Cognitive Services; roles for Microsoft solution accelerators and managed solutions; and integration of the Azure footprint with legacy infrastructure. The book concludes with a discussion of best practices in defining and developing solutions and creating a plan for success. What You Will Learn Design the right IoT architecture to deliver solutions for a variety of project needs Connect IoT devices to Azure for data collection and delivery of services Use Azure Machine Learning and Cognitive Services to deliver intelligence in cloud-based solutions and at the edge Understand the benefits and tradeoffs of Microsoft's solution accelerators and managed solutions Investigate new use cases that are described and apply best practices in deployment strategies Integrate cutting-edge Azure deployments with existing legacy data sources Who This Book Is For Developers and architects new to IoT projects or new to Microsoft Azure IoT components as well as readers interested in best practices used in architecting IoT solutions that utilize the Azure platform

A handy reference guide for data analysts and data scientists to help to obtain value from big data analytics using Spark on Hadoop clusters About This Book This book is based on the latest 2.0 version of Apache Spark and 2.7 version of Hadoop integrated with most commonly used tools. Learn all Spark stack components including latest topics such as DataFrames, DataSets, GraphFrames, Structured Streaming, DataFrame based ML Pipelines and SparkR. Integrations with frameworks such as

HDFS, YARN and tools such as Jupyter, Zeppelin, NiFi, Mahout, HBase Spark Connector, GraphFrames, H2O and Hivemall. Who This Book Is For Though this book is primarily aimed at data analysts and data scientists, it will also help architects, programmers, and practitioners. Knowledge of either Spark or Hadoop would be beneficial. It is assumed that you have basic programming background in Scala, Python, SQL, or R programming with basic Linux experience. Working experience within big data environments is not mandatory. What You Will Learn Find out and implement the tools and techniques of big data analytics using Spark on Hadoop clusters with wide variety of tools used with Spark and Hadoop Understand all the Hadoop and Spark ecosystem components Get to know all the Spark components: Spark Core, Spark SQL, DataFrames, DataSets, Conventional and Structured Streaming, MLLib, ML Pipelines and Graphx See batch and real-time data analytics using Spark Core, Spark SQL, and Conventional and Structured Streaming Get to grips with data science and machine learning using MLLib, ML Pipelines, H2O, Hivemall, Graphx, SparkR and Hivemall. In Detail Big Data Analytics book aims at providing the fundamentals of Apache Spark and Hadoop. All Spark components – Spark Core, Spark SQL, DataFrames, Data sets, Conventional Streaming, Structured Streaming, MLlib, Graphx and

Hadoop core components – HDFS, MapReduce and Yarn are explored in greater depth with implementation examples on Spark + Hadoop clusters. It is moving away from MapReduce to Spark. So, advantages of Spark over MapReduce are explained at great depth to reap benefits of in-memory speeds. DataFrames API, Data Sources API and new Data set API are explained for building Big Data analytical applications. Real-time data analytics using Spark Streaming with Apache Kafka and HBase is covered to help building streaming applications. New Structured streaming concept is explained with an IOT (Internet of Things) use case. Machine learning techniques are covered using MLLib, ML Pipelines and SparkR and Graph Analytics are covered with GraphX and GraphFrames components of Spark. Readers will also get an opportunity to get started with web based notebooks such as Jupyter, Apache Zeppelin and data flow tool Apache NiFi to analyze and visualize data. Style and approach This step-by-step pragmatic guide will make life easy no matter what your level of experience. You will deep dive into Apache Spark on Hadoop clusters through ample exciting real-life examples. Practical tutorial explains data science in simple terms to help programmers and data analysts get started with Data Science

DESIGNING BIG DATA PLATFORMS Provides expert guidance and valuable insights on getting the most out of Big Data systems An array of tools are currently available for managing and processing data—some are ready-to-go solutions that can be immediately deployed, while others require complex and time-intensive setups. With such a vast range of options, choosing the right tool to build a solution can be complicated, as can determining which tools work well with each other. Designing Big Data Platforms provides clear and authoritative guidance on the critical decisions necessary for successfully deploying, operating, and maintaining Big Data systems. This highly practical guide helps readers understand how to process large amounts of data with well-known Linux tools and database solutions, use effective techniques to collect and manage data from multiple sources, transform data into meaningful business insights, and much more. Author Yusuf Aytas, a software engineer with a vast amount of big data experience, discusses the design of the ideal Big Data platform: one that meets the needs of data analysts, data engineers, data scientists, software engineers, and a spectrum of other stakeholders across an organization. Detailed yet accessible chapters cover key topics such as stream data processing, data analytics, data science, data discovery, and data security. This real-world manual for Big Data technologies: Provides up-to-date coverage of the tools currently used in Big Data processing and management Offers step-by-step guidance on building a data pipeline, from basic scripting to distributed systems Highlights and explains how data is processed at scale Includes an introduction to the foundation of a modern data platform Designing Big Data Platforms: How to Use, Deploy, and Maintain Big Data Systems is a must-have for all professionals working with Big Data, as well researchers and students in computer science and related fields.

?? ??, ??? ??? ??? ????? ????. ??? ??? ??? ?????? ?? ??? ??? ?? ???? ???. ?? ??? ??, ?? ??? ???? ???? ???? ??????, ???? ????, ?? ???, ?? ?? ??, ??? ?? ??, ?????? ????, ??? ????? ???. ? ????? ?? ??? ???? ??? ??? ???? ?? ???? ???? ??? ??(???, ??? XD, ??, ??, ??, ????, ???, ???, R, Rjava)? ???? ?????. ?? ?? ????? ?? ??? ? ? ??? ????? ? ??? ??? ??? ?? ? ??.

This book constitutes the thoroughly refereed post-conference proceedings of the 12th International Meeting on Computational Intelligence Methods for Bioinformatics and Biostatistics, CIBB 2015, held in Naples, Italy, in September, 2015. The 21 revised full papers presented were carefully reviewed and selected from 24 submissions. They present problems concerning computational techniques in bioinformatics, systems biology and medical informatics discussing cutting edge methodologies and accelerate life science discoveries, as well as novel challenges with an high impact on molecular biology and translational medicine.

“This book is a critically needed resource for the newly released Apache Hadoop 2.0, highlighting YARN as the significant breakthrough that broadens Hadoop beyond the MapReduce paradigm.” —From the Foreword by Raymie Stata, CEO of Altiscale The Insider’s Guide to Building Distributed, Big Data Applications with Apache Hadoop™ YARN Apache Hadoop is helping drive the Big Data revolution. Now, its data processing has been completely overhauled: Apache Hadoop YARN provides resource management at data center scale and easier ways to create distributed applications that process petabytes of data. And now in Apache Hadoop™ YARN, two Hadoop technical leaders show you how to develop new applications and adapt existing code to fully leverage these revolutionary advances. YARN project founder Arun Murthy and project lead Vinod Kumar Vavilapalli demonstrate how YARN increases scalability and cluster utilization, enables new programming models and services, and opens new options beyond Java and batch processing. They walk you through the entire YARN project lifecycle, from installation through deployment. You’ll find many examples drawn from the authors’ cutting-edge experience—first as Hadoop’s earliest developers and implementers at Yahoo! and now as Hortonworks developers moving the platform forward and helping customers succeed with it. Coverage includes YARN’s goals, design, architecture, and components—how it expands the Apache Hadoop ecosystem Exploring YARN on a single node Administering YARN clusters and Capacity Scheduler Running existing MapReduce applications Developing a large-scale clustered YARN application Discovering new open source frameworks that run under YARN

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

Learn all you need to know about seven key innovations disrupting business analytics today. These innovations—the open source business model, cloud analytics, the Hadoop ecosystem, Spark and in-memory analytics, streaming analytics, Deep Learning, and self-service analytics—are radically changing how businesses use data for competitive advantage. Taken together, they are disrupting the business analytics value chain, creating new opportunities. Enterprises who seize the opportunity will thrive and prosper, while others struggle and decline: disrupt or be disrupted. Disruptive Business Analytics provides strategies to profit from disruption. It shows you how to organize for insight, build and provision an open source stack, how to practice lean data warehousing, and how to assimilate disruptive innovations into an organization. Through a short history of business analytics and a detailed survey of products and services, analytics authority Thomas W. Dinsmore provides a practical explanation of the most compelling innovations available today. What You'll Learn Discover how the open source business model works and how to make it work for you See how cloud computing completely changes the economics of analytics Harness the power of Hadoop and its ecosystem Find out why Apache Spark is everywhere Discover the potential of streaming and real-time analytics Learn what Deep Learning can do and why it matters See how self-service analytics can change the way organizations do business Who This Book Is For Corporate actors at all levels of responsibility for analytics: analysts, CIOs, CTOs, strategic decision makers, managers, systems architects, technical marketers, product developers, IT personnel, and consultants.

If you have a working knowledge of Hadoop 1.x but want to start afresh with YARN, this book is ideal for you. You will be able to install and administer a YARN cluster and also discover the configuration settings to fine-tune your cluster both in terms of performance and scalability. This book will help you develop, deploy, and run multiple applications/frameworks on the same shared YARN cluster.

This book constitutes the refereed proceedings of the 5th International Symposium on Cyber Security Cryptography and Machine Learning, CSCML 2021, held in Be'er Sheva, Israel, in July 2021. The 22 full and 13 short papers presented together with a keynote paper in this volume were carefully reviewed and selected from 48 submissions. They deal with the theory, design, analysis, implementation, or application of cyber security, cryptography and machine learning systems and networks, and conceptually innovative topics in these research areas.

"Hadoop and Spark Fundamentals LiveLessons provides 9+ hours of video introduction to the Apache Hadoop Big Data ecosystem. The tutorial includes background information and explains the core components of Hadoop, including Hadoop Distributed File Systems (HDFS), MapReduce, the YARN resource manager, and YARN Frameworks. In addition, it demonstrates how to use Hadoop at several levels, including the native Java interface, C++ pipes, and the universal streaming program interface. Examples include how to use benchmarks and high-level tools, including the Apache Pig scripting language, Apache Hive "SQL-like" interface, Apache Flume for streaming input, Apache Sqoop for import and export of relational data, and Apache Oozie for Hadoop workflow management. In addition, there is comprehensive coverage of Spark, PySpark, and the Zeppelin web-GUI. The steps for easily installing a working Hadoop/Spark system on a desktop/laptop and on a local stand-alone cluster using the powerful Ambari GUI are also included. All software used in these LiveLessons is open source and freely available for your use and experimentation. A bonus lesson includes a quick primer on the Linux command line as used with Hadoop and Spark."--Resource description page.

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. This book covers three major parts of Big Data: concepts, theories and applications. Written by world-renowned leaders in Big Data, this book explores the problems, possible solutions and directions for Big Data in research and practice. It also focuses on high level concepts such as definitions of Big Data from different angles; surveys in research and applications; and existing tools, mechanisms, and systems in practice. Each chapter is independent from the other chapters, allowing users to read any chapter directly. After examining the practical side of Big Data, this book presents theoretical perspectives. The theoretical research ranges from Big Data representation, modeling and topology to distribution and dimension reducing. Chapters also investigate the many disciplines that involve Big Data, such as statistics, data mining, machine learning, networking, algorithms, security and differential geometry. The last section of this book introduces Big Data applications from different communities, such as business, engineering and science. Big Data Concepts, Theories and Applications is designed as a reference for researchers and advanced level students in computer science, electrical engineering and mathematics. Practitioners who focus on information systems, big data, data mining, business analysis and other related fields will also find this material valuable.