

## Getting Started With Orientdb

Run blazingly fast queries on complex graph datasets with the power of the Neo4j graph database About This Book Get acquainted with graph database systems and apply them in real-world use cases Use Cypher query language, APOC and other Neo4j extensions to derive meaningful analysis from complex data sets. A practical guide filled with ready to use examples on querying, graph processing and visualizing information to build smarter spatial applications. Who This Book Is For This book is for developers who want an alternative way to store and process data within their applications. No previous graph database experience is required; however, some basic database knowledge will help you understand the concepts more easily. What You Will Learn Understand the science of graph theory, databases and its advantages over traditional databases. Install Neo4j, model data and learn the most common practices of traversing data Learn the Cypher query language and tailor-made procedures to analyze and derive meaningful representations of data Improve graph techniques with the help of precise procedures in the APOC library Use Neo4j advanced extensions and plugins for performance optimization. Understand how Neo4j's new security features and clustering architecture are used for large scale deployments. In Detail Neo4j is a graph database that allows traversing huge amounts of data with ease. This book aims at quickly getting you started with the popular graph database Neo4j. Starting with a brief introduction to graph theory, this book will show you the advantages of using graph databases along with data modeling techniques for graph databases. You'll gain practical hands-on experience with commonly used and lesser known features for updating graph store with Neo4j's Cypher query language. Furthermore, you'll also learn to create awesome procedures using APOC and extend Neo4j's functionality, enabling integration, algorithmic analysis, and other advanced spatial operation capabilities on data. Through the course of the book you will come across implementation examples on the latest updates in Neo4j, such as in-graph indexes, scaling, performance improvements, visualization, data refactoring techniques, security enhancements, and much more. By the end of the book, you'll have gained the skills to design and implement modern spatial applications, from graphing data to unraveling business capabilities with the help of real-world use cases. Style and approach A step-by-step approach of adopting Neo4j, the world's leading graph database. This book includes a lot of background information, helps you grasp the fundamental concepts behind this radical new way of dealing with connected data, and will give you lots of examples of use cases and environments where a graph database would be a great fit

If you are a professional or enthusiast who has a basic understanding of graphs or has basic knowledge of Neo4j operations, this is the book for you. Although it is targeted at an advanced user base, this book can be used by beginners as it touches upon the basics. So, if you are passionate about taming complex data with the help of graphs and building high performance applications, you will be able to get valuable insights from this book.

Congratulations! You completed the MongoDB application within the given tight timeframe and there is a party to celebrate your application's release into production. Although people are congratulating you at the celebration, you are feeling some uneasiness inside. To complete the project on time required making a lot of assumptions about the data, such as what terms meant and how calculations are derived. In addition, the poor documentation about the application will be of limited use to the support team, and not investigating all of the inherent rules in the data may eventually lead to poorly-performing structures in the not-so-distant future. Now, what if you had a time machine and could go back and read this book. You would learn that even NoSQL databases like MongoDB require some level of data modeling. Data modeling is the process of learning about the data, and regardless of technology, this process must be performed for a successful application. You would learn the value of conceptual, logical, and physical data modeling and how each stage increases our knowledge of the data and reduces assumptions and poor design decisions. Read this book to learn how to do data modeling for MongoDB applications, and accomplish these five objectives: Understand how data modeling contributes to the process of learning about the data, and is, therefore, a required technique, even when the resulting database is not relational. That is, NoSQL does not mean NoDataModeling! Know how NoSQL databases differ from traditional relational databases, and where MongoDB fits. Explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts, and learn the basics of adding, querying, updating, and deleting data in MongoDB. Practice a streamlined, template-driven approach to performing conceptual, logical, and physical data modeling. Recognize that data modeling does not always have to lead to traditional data models! Distinguish top-down from bottom-up development approaches and complete a top-down case study which ties all of the modeling techniques together. This book is written for anyone who is working with, or will be working with MongoDB, including business analysts, data modelers, database administrators, developers, project managers, and data scientists. There are three sections: In Section I, Getting Started, we will reveal the power of data modeling and the tight connections to data models that exist when designing any type of database (Chapter 1), compare NoSQL with traditional relational databases and where MongoDB fits (Chapter 2), explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts (Chapter 3), and explain the basics of adding, querying, updating, and deleting data in MongoDB (Chapter 4). In Section II, Levels of Granularity, we cover Conceptual Data Modeling (Chapter 5), Logical Data Modeling (Chapter 6), and Physical Data Modeling (Chapter 7). Notice the "ing" at the end of each of these chapters. We focus on the process of building each of these models, which is where we gain essential business knowledge. In Section III, Case Study, we will explain both top down and bottom up development approaches and go through a top down case study where we start with business requirements and end with the MongoDB database. This case study will tie together all of the techniques in the previous seven chapters. Nike Senior Data Architect Ryan Smith wrote the foreword. Key points are included at the end of each chapter as a way to reinforce concepts. In addition, this book is loaded with hands-on exercises, along with their answers provided in Appendix A. Appendix B contains all of the book's references and Appendix C contains a glossary of the terms used throughout the text.

Big Data Analytics is on the rise in the last years of the current decade. Data are overwhelming the computation capacity of high performance servers. Cloud, grid, edge and fog computing are a few examples of the current hype. Computational Intelligence offers two faces to deal with the development of models: on the one hand, the crisp approach, which considers for every variable an exact value and, on the other hand, the fuzzy focus, which copes with values between two boundaries. This book presents 114 papers from the 4th International Conference on Fuzzy Systems and Data Mining (FSDM 2018), held in Bangkok, Thailand, from 16 to 19 November 2018. All papers were carefully reviewed by program committee members, who took into consideration the breadth and depth of the research topics that fall within the scope of FSDM. The acceptance rate was 32.85% . Offering a state-of-the-art overview of fuzzy systems and data mining, the publication will be of interest to all those whose work involves data science. This book is for developers who want an alternative way to store and process data within their applications. No previous graph database experience is required; however, some basic database knowledge will help you understand the concepts more easily.

This book constitutes the thoroughly refereed post-conference proceedings of the First International Workshop on Algorithmic Aspects of Cloud Computing, ALGO CLOUD 2015, held in Patras, Greece, in September 2015 in conjunction with ALGO 2015. The 13 revised full papers presented together with 2 tutorial papers were carefully reviewed and selected from 37 initial submissions. They cover a wide range of topics in two main tracks: algorithmic aspects of large-scale data stores, and software tools and distributed architectures for cloud-based data management.

NoSQL was developed to overcome the limitations of relational databases in the largest Web applications at companies such as Google, Yahoo and Facebook. As it is applied more widely, developers are finding that it can simplify scalability while requiring far less coding and management overhead. However, NoSQL requires fundamentally different approaches to database design and modeling, and many conventional relational techniques lead to suboptimal results. *¿ NoSQL for Mere Mortals* is an easy, practical guide to succeeding with NoSQL in your environment. Following the classic, best-selling format pioneered in *SQL Queries for Mere Mortals*, enterprise database expert Dan Sullivan guides you step-by-step through choosing technologies, designing high-performance databases, and planning for long-term maintenance. *¿ Sullivan* introduces each type of NoSQL database, shows how to install and manage them, and demonstrates how to leverage their features while avoiding common mistakes that lead to poor performance and unmet requirements. He uses four popular NoSQL databases as reference models: MongoDB, a document database; Cassandra, a column family data store; Redis, a key-value database; and Neo4j, a graph database. You'll find explanations of each database's structure and capabilities, practical guidelines for choosing amongst them, and expert guidance on designing databases with them. *¿ Packed with examples, NoSQL for Mere Mortals* is today's best way to master NoSQL -- whether you're a DBA, developer, user, or student.

This book provides a review of advanced topics relating to the theory, research, analysis and implementation in the context of big data platforms and their applications, with a focus on methods, techniques, and performance evaluation. The explosive growth in the volume, speed, and variety of data being produced every day requires a continuous increase in the processing speeds of servers and of entire network infrastructures, as well as new resource management models. This poses significant challenges (and provides striking development opportunities) for data intensive and high-performance computing, i.e., how to efficiently turn extremely large datasets into valuable information and meaningful knowledge. The task of context data management is further complicated by the variety of sources such data derives from, resulting in different data formats, with varying storage, transformation, delivery, and archiving requirements. At the same time rapid responses are needed for real-time applications. With the emergence of cloud infrastructures, achieving highly scalable data management in such contexts is a critical problem, as the overall application performance is highly dependent on the properties of the data management service.

This book will introduce Redis and help you understand its various facets. Starting with an introduction to NoSQL, you will learn how to install Redis and how to classify and work with data structures. By working with real world scenarios pertaining to using Redis, you will discover sharding and indexing techniques, along with how to improve scalability and performance through persistent strategies and data migration techniques. With the help of multiple examples, you will learn to design web and business applications. You will also learn how to configure Redis for setting up clusters and tuning it for performance. At the end of this book, you will find essential tips on backup and recovery strategies for the Redis environment.

to date, research on interactive intelligent systems has largely focused either on the realisation of the systems' capabilities or on the cognitive processes and/or behaviour of their users. With the rapid development of Internet-based technologies, the design of interactive intelligent systems is facing many emerging issues and challenges such as investigating the ways that artificial agents and human intelligence can collaborate for better performance, understanding user requirements and user cognitive processes, safeguarding user privacy, etc.

This book provides the latest research findings and developments in the field of interactive intelligent systems, addressing diverse areas such as autonomous systems, Internet and cloud computing, pattern recognition and vision systems, mobile computing and intelligent networking, and e-enabled systems. It gathers selected papers from the International Conference on Intelligent and Interactive Systems and Applications (IISA2016) held on June 25–26, 2016 in Shanghai, China. Interactive intelligent systems are among the most important multi-disciplinary research and development domains of artificial intelligence, human–computer interaction, machine learning and new Internet-based technologies. Accordingly, these systems embrace a considerable number of application areas such as autonomous systems, expert systems, mobile systems, recommender systems, knowledge-based and semantic web-based systems, virtual communication environments, and decision support systems, to name a few. To date, research on interactive intelligent systems has largely focused either on the realisation of the systems' capabilities or on the cognitive processes and/or behaviour of their users. With the rapid development of Internet-based technologies, the design of interactive intelligent systems is facing many emerging issues and challenges such as investigating the ways that artificial agents and human intelligence can collaborate for better performance, understanding user requirements and user cognitive processes, safeguarding user privacy, etc.

Data is getting bigger and more complex by the day, and so are your choices in handling it. Explore some of the most cutting-edge databases available - from a traditional relational database to newer NoSQL approaches - and make informed decisions about challenging data storage problems. This is the only comprehensive guide to the world of NoSQL databases, with in-depth practical and conceptual introductions to seven different technologies: Redis, Neo4J, CouchDB, MongoDB, HBase, Postgres, and DynamoDB. This second edition includes a new chapter on DynamoDB and updated content for each chapter. While relational databases such as MySQL remain as relevant as ever, the alternative, NoSQL paradigm has opened up new horizons in performance and scalability and changed the way we approach data-centric problems. This book presents the essential concepts behind each database alongside hands-on examples that make each technology come alive. With each database, tackle a real-world problem that highlights the concepts and features that make it shine. Along the way, explore five database models - relational, key/value, columnar, document, and graph - from the perspective of challenges faced by real applications. Learn how MongoDB and CouchDB are strikingly different, make your applications faster with Redis and more connected with Neo4J, build a cluster of HBase servers using cloud services such as Amazon's Elastic MapReduce, and more. This new edition brings a brand new chapter on DynamoDB, updated code samples and exercises, and a more up-to-date account of each database's feature set. Whether you're a programmer building the next big thing, a data scientist seeking solutions to thorny problems, or a technology enthusiast venturing into new territory, you will find something to inspire you in this book. What You Need: You'll need a \*nix shell (Mac OS or Linux preferred, Windows users will need Cygwin), Java 6 (or greater), and Ruby 1.8.7 (or greater). Each chapter will list the downloads required for that database. Big Data Systems encompass massive challenges related to data diversity, storage mechanisms, and requirements of massive computational power. Further, capabilities of big

data systems also vary with respect to type of problems. For instance, distributed memory systems are not recommended for iterative algorithms. Similarly, variations in big data systems also exist related to consistency and fault tolerance. The purpose of this book is to provide a detailed explanation of big data systems. The book covers various topics including Networking, Security, Privacy, Storage, Computation, Cloud Computing, NoSQL and NewSQL systems, High Performance Computing, and Deep Learning. An illustrative and practical approach has been adopted in which theoretical topics have been aided by well-explained programming and illustrative examples. Key Features: Introduces concepts and evolution of Big Data technology. Illustrates examples for thorough understanding. Contains programming examples for hands on development. Explains a variety of topics including NoSQL Systems, NewSQL systems, Security, Privacy, Networking, Cloud, High Performance Computing, and Deep Learning. Exemplifies widely used big data technologies such as Hadoop and Spark. Includes discussion on case studies and open issues. Provides end of chapter questions for enhanced learning. A standard tutorial aimed at making you an OrientDB expert, through the use of practical examples, explained in a step-by-step format. Getting Started with OrientDB 1.3.0 is great for database designers, developers, and systems engineers. It is assumed that you are familiar with NoSQL concepts, Java, and networking principles.

The recent discovery of small and long non-coding RNAs (ncRNAs) has represented a major breakthrough in the life sciences. These molecules add a new layer of complexity to biological processes and pathways by revealing a sophisticated and dynamic interconnected system whose structure is just beginning to be uncovered. Genetic and epigenetic aberrations affecting ncRNA gene sequences and their expression have been linked to a variety of pathological conditions, including cancer, cardiovascular and neurological diseases. Latest advances in the development of high throughput analysis techniques may help to shed light on the complex regulatory mechanisms in which ncRNA molecules are involved. Bioinformatics tools constitute a unique and essential resource for non-coding RNA studies, providing a powerful technology to organize, integrate and analyze the huge amount of data produced daily by wet biology experiments in order to discover patterns, identify relationships among heterogeneous biological elements and formulate functional hypotheses. This Research Topic reviews current knowledge, introduces novel methods, and discusses open challenges of this exciting and innovative field in connection with the most important biomedical applications. It consists of four reviews and six original research and methods articles, spanning the full scope of the Research Topic.

Hibernate and MongoDB are a powerful combination of open source persistence and NoSQL technologies for today's Java-based enterprise and cloud application developers. Hibernate is the leading open source Java-based persistence, object relational management engine, recently repositioned as an object grid management engine. MongoDB is a growing, popular open source NoSQL framework, especially popular among cloud application and big data developers. With these two, enterprise and cloud developers have a "complete out of the box" solution. Pro Hibernate and MongoDB shows you how to use and integrate Hibernate and MongoDB. More specifically, this book guides you through the bootstrap; building transactions; handling queries and query entities; and mappings. Then, this book explores the principles and techniques for taking these application principles to the cloud, using the OpenShift Platform as a Service (PaaS) and more. In this book, you get two case studies: An enterprise application using Hibernate and MongoDB. then, A cloud application (OpenShip) migrated from the enterprise application case study After reading or using this book, you come away with the experience from two case studies that give you possible frameworks or templates that you can apply to your own specific application or cloud application building context.

Persistence in PHP with the Doctrine ORM is a concise, fast, and focused guide to build a blog engine with advanced features such as native queries and lifecycle callbacks. This book is primarily intended for PHP developers and architects who want to increase their skills in the field of Persistence and ORM to map the data they are working on to objects they are using in programming. Basic knowledge of databases and PDO and working knowledge of PHP namespaces is a prerequisite.

Leverage existing free open source software to build an identity and access management (IAM) platform that can serve your organization for the long term. With the emergence of open standards and open source software, it's now easier than ever to build and operate your own IAM stack. The most common culprit of the largest hacks has been bad personal identification. In terms of bang for your buck, effective access control is the best investment you can make. Financially, it's more valuable to prevent than to detect a security breach. That's why Identity and Access Management (IAM) is a critical component of an organization's security infrastructure. In the past, IAM software has been available only from large enterprise software vendors. Commercial IAM offerings are bundled as "suites" because IAM is not just one component. It's a number of components working together, including web, authentication, authorization, cryptographic, and persistence services. Securing the Perimeter documents a recipe to take advantage of open standards to build an enterprise-class IAM service using free open source software. This recipe can be adapted to meet the needs of both small and large organizations. While not a comprehensive guide for every application, this book provides the key concepts and patterns to help administrators and developers leverage a central security infrastructure. Cloud IAM service providers would have you believe that managing an IAM is too hard. Anything unfamiliar is hard, but with the right road map, it can be mastered. You may find SaaS identity solutions too rigid or too expensive. Or perhaps you don't like the idea of a third party holding the credentials of your users—the keys to your kingdom. Open source IAM provides an alternative. Take control of your IAM infrastructure if digital services are key to your organization's success. What You'll Learn Understand why you should deploy a centralized authentication and policy management infrastructure Use the SAML or Open ID Standards for web or single sign-on, and OAuth for API Access Management Synchronize data from existing identity repositories such as Active Directory Deploy two-factor authentication services Who This Book Is For Security architects (CISO, CSO), system engineers/administrators, and software developers

Summary Introducing Data Science teaches you how to accomplish the fundamental tasks that occupy data scientists. Using the Python language and common Python libraries, you'll experience firsthand the challenges of dealing with data at scale and gain a solid foundation in data science. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Many companies need developers with data science skills to work on projects ranging from social media marketing to machine learning. Discovering what you need to learn to begin a career as a data scientist can seem bewildering. This book is designed to help you get started. About the Book Introducing Data Science Introducing Data Science explains vital data science concepts and teaches you how to accomplish the fundamental tasks that occupy data scientists. You'll explore data visualization, graph databases, the use of NoSQL, and the data science process. You'll use the Python language and common Python libraries as you experience firsthand the challenges of dealing with data at scale. Discover how Python allows you to gain insights from data sets so big that they need to be stored on multiple machines, or from data moving so quickly that no single machine can handle it. This book gives you hands-on experience with the most popular Python data science libraries, Scikit-learn and StatsModels. After reading this book, you'll have the solid foundation you need to start a career in data science. What's Inside Handling large data Introduction to machine learning Using Python to work with data Writing data science algorithms About the Reader This book assumes you're comfortable reading code in Python or a similar language, such as C, Ruby, or JavaScript. No prior experience with data science is required. About the Authors Davy Cielen, Arno D. B. Meysman, and Mohamed Ali are the founders and managing partners of Optimately and Maiton, where they focus on developing data science projects and solutions in various sectors. Table of Contents Data science in a big data world The data science process Machine learning Handling large data on a single computer First steps in big data Join the NoSQL movement The rise of graph databases Text mining and text analytics Data visualization to the end user

Ubuntu Unleashed 2021 Edition is filled with unique information for everyone who wants to make the most of the Ubuntu Linux operating system, including the latest in Ubuntu mobile development. This new edition has been thoroughly updated by a long-time Ubuntu community leader to reflect the new Ubuntu 20.04 and the forthcoming Ubuntu 20.10 and 21.04 releases. Former Ubuntu Forum administrator Matthew Helmke covers all you need to know about Ubuntu 20.04 installation, configuration, productivity, multimedia, development, system administration, server operations, networking, virtualization, security, DevOps, and more - including intermediate-to-advanced techniques you won't find in any other book. Helmke introduces Ubuntu's key productivity and Web development tools, programming languages, hardware support, and more. You'll find new and improved coverage of Ubuntu's Unity interface, various types of servers, software repositories, database options, virtualization and cloud services, development tools, monitoring, troubleshooting, Ubuntu's push into mobile and other touch screen devices, and much more. The companion DVD includes the full Ubuntu 20.04 distribution as well as the complete LibreOffice office suite and hundreds of additional programs and utilities. Ubuntu Unleashed provides detailed information on how to... Configure and customize the Unity desktop Get started with multimedia and productivity applications, including LibreOffice Manage Linux services, users, and software packages Administer and run Ubuntu from the command line Automate tasks and use shell scripting Provide secure remote access and configure a secure VPN Manage kernels and modules Administer file, print, email, proxy, LDAP, DNS, and HTTP servers (Apache, Nginx, or alternatives) Learn about new options for managing large numbers of servers Work with databases (both SQL and the newest NoSQL alternatives) Get started with virtualization Build a private cloud with Juju and Charms Learn the basics about popular programming languages including Python, PHP, Perl, and new alternatives such as Go and Rust Learn about Ubuntu's work toward usability on touch-screen and phone devices

This book offers a comprehensive introduction to relational (SQL) and non-relational (NoSQL) databases. The authors thoroughly review the current state of database tools and techniques, and examine coming innovations. The book opens with a broad look at data management, including an overview of information systems and databases, and an explanation of contemporary database types: SQL and NoSQL databases, and their respective management systems The nature and uses of Big Data A high-level view of the organization of data management Data Modeling and Consistency Chapter-length treatment is afforded Data Modeling in both relational and graph databases, including enterprise-wide data architecture, and formulas for database design. Coverage of languages extends from an overview of operators, to SQL and and QBE (Query by Example), to integrity constraints and more. A full chapter probes the challenges of Ensuring Data Consistency, covering: Multi-User Operation Troubleshooting Consistency in Massive Distributed Data Comparison of the ACID and BASE consistency models, and more System Architecture also gets from its own chapter, which explores Processing of Homogeneous and Heterogeneous Data; Storage and Access Structures; Multi-dimensional Data Structures and Parallel Processing with MapReduce, among other topics. Post-Relational and NoSQL Databases The chapter on post-relational databases discusses the limits of SQL – and what lies beyond, including Multi-Dimensional Databases, Knowledge Bases and and Fuzzy Databases. A final chapter covers NoSQL Databases, along with Development of Non-Relational Technologies, Key-Value, Column-Family and Document Stores XML Databases and Graphic Databases, and more The book includes more than 100 tables, examples and illustrations, and each chapter offers a list of resources for further reading. SQL & NoSQL Databases conveys the strengths and weaknesses of relational and non-relational approaches, and shows how to undertake development for big data applications. The book benefits readers including students and practitioners working across the broad field of applied information technology. This textbook has been recommended and developed for university courses in Germany, Austria and Switzerland.

Develop your red team skills by learning essential foundational tactics, techniques, and procedures, and boost the overall security posture of your organization by leveraging the homefield advantage Key Features Build, manage, and measure an offensive red team program Leverage the homefield advantage to stay ahead of your adversaries Understand core adversarial tactics and techniques, and protect pentesters and pentesting assets Book Description It's now more important than ever for organizations to be ready to detect and respond to security events and breaches. Preventive measures alone are not enough for dealing with adversaries. A well-rounded prevention, detection, and response program is required. This book will guide you through the stages of building a red team program, including strategies and homefield advantage opportunities to boost security. The book starts by guiding you through establishing, managing, and measuring a red team program, including effective ways for sharing results and findings to raise awareness. Gradually, you'll learn about progressive operations such as cryptocurrency mining, focused privacy testing, targeting telemetry, and even blue team tooling. Later, you'll discover knowledge graphs and how to build them, then become well-versed with basic to advanced techniques related to hunting for credentials, and learn to automate Microsoft Office and browsers to your advantage. Finally, you'll get to grips with protecting assets using decoys, auditing, and alerting with examples for major operating systems. By the end of this book, you'll have learned how to build, manage, and measure a red team program effectively and be well-versed with the fundamental operational techniques required to enhance your existing skills. What you will learn Understand the risks associated with security breaches Implement strategies for building an effective penetration testing team Map out the homefield using knowledge graphs Hunt credentials using indexing and other practical techniques Gain blue team tooling insights to enhance your red team skills Communicate results and influence decision makers with appropriate data Who this book is for This is one of the few detailed cybersecurity books for

penetration testers, cybersecurity analysts, security leaders and strategists, as well as red team members and chief information security officers (CISOs) looking to secure their organizations from adversaries. The program management part of this book will also be useful for beginners in the cybersecurity domain. To get the most out of this book, some penetration testing experience, and software engineering and debugging skills are necessary.

Discover how graph algorithms can help you leverage the relationships within your data to develop more intelligent solutions and enhance your machine learning models. You'll learn how graph analytics are uniquely suited to unfold complex structures and reveal difficult-to-find patterns lurking in your data. Whether you are trying to build dynamic network models or forecast real-world behavior, this book illustrates how graph algorithms deliver value—from finding vulnerabilities and bottlenecks to detecting communities and improving machine learning predictions. This practical book walks you through hands-on examples of how to use graph algorithms in Apache Spark and Neo4j—two of the most common choices for graph analytics. Also included: sample code and tips for over 20 practical graph algorithms that cover optimal pathfinding, importance through centrality, and community detection. Learn how graph analytics vary from conventional statistical analysis Understand how classic graph algorithms work, and how they are applied Get guidance on which algorithms to use for different types of questions Explore algorithm examples with working code and sample datasets from Spark and Neo4j See how connected feature extraction can increase machine learning accuracy and precision Walk through creating an ML workflow for link prediction combining Neo4j and Spark

The need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational “NoSQL” databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage's pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Polyglot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access.

Graph Databases in Action teaches readers everything they need to know to begin building and running applications powered by graph databases. Right off the bat, seasoned graph database experts introduce readers to just enough graph theory, the graph database ecosystem, and a variety of datastores. They also explore modelling basics in action with real-world examples, then go hands-on with querying, coding traversals, parsing results, and other essential tasks as readers build their own graph-backed social network app complete with a recommendation engine! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

This book constitutes the refereed post-conference proceedings of the 6th TPC Technology Conference, TPCTC 2014, held in Hangzhou, China, in September 2014. It contains 12 selected peer-reviewed papers, a report from the TPC Public Relations Committee. Many buyers use TPC benchmark results as points of comparison when purchasing new computing systems. The information technology landscape is evolving at a rapid pace, challenging industry experts and researchers to develop innovative techniques for evaluation, measurement and characterization of complex systems. The TPC remains committed to developing new benchmark standards to keep pace and one vehicle for achieving this objective is the sponsorship of the Technology Conference on Performance Evaluation and Benchmarking (TPCTC). Over the last five years TPCTC has been held successfully in conjunction with VLDB.

The topic of NoSQL databases has recently emerged, to face the Big Data challenge, namely the ever increasing volume of data to be handled. It is now recognized that relational databases are not appropriate in this context, implying that new database models and techniques are needed. This book presents recent research works, covering the following basic aspects: semantic data management, graph databases, and big data management in cloud environments. The chapters in this book report on research about the evolution of basic concepts such as data models, query languages, and new challenges regarding implementation issues.

This book constitutes the thoroughly refereed short papers, workshops and doctoral consortium papers of the 23rd European Conference on Advances in Databases and Information Systems, ADBIS 2019, held in Bled, Slovenia, in September 2019. The 19 short research papers and the 5 doctoral consortium papers were carefully reviewed and selected from 103 submissions, and the 31 workshop papers were selected out of 67 submitted papers. The papers are organized in the following sections: Short Papers; Workshops Papers; Doctoral Consortium Papers; and cover a wide spectrum of topics related to database and information systems technologies for advanced applications.

Get up to speed on the nuances of NoSQL databases and what they mean for your organization This easy to read guide to NoSQL databases provides the type of no-nonsense overview and analysis that you need to learn, including what NoSQL is and which database is right for you. Featuring specific evaluation criteria for NoSQL databases, along with a look into the pros and cons of the most popular options, NoSQL For Dummies provides the fastest and easiest way to dive into the details of this incredible technology. You'll gain an understanding of how to use NoSQL databases for mission-critical enterprise architectures and projects, and real-world examples reinforce the primary points to create an action-oriented resource for IT pros. If you're planning a big data project or platform, you probably already know you need to select a NoSQL database to complete your architecture. But with options flooding the market and updates and add-ons coming at a rapid pace, determining what you require now, and in the future, can be a tall task. This is where NoSQL For Dummies comes in! Learn the basic tenets of NoSQL databases and why they have come to the forefront as data has outpaced the capabilities of relational databases Discover major players among NoSQL databases, including Cassandra, MongoDB, MarkLogic, Neo4J, and others Get an in-depth look at the benefits and disadvantages of the wide variety of NoSQL database options Explore the needs of your organization as they relate to the capabilities of specific NoSQL databases Big data and Hadoop get all the attention, but when it comes down to it, NoSQL databases are the engines that power many big data analytics initiatives. With NoSQL For Dummies, you'll go beyond relational databases to ramp up your enterprise's data architecture in no time.

If you are already using Neo4j in your application and want to learn more about data analysis or database graphs, this is the book for you. This book also caters for your needs if you are looking to migrate your existing application to Neo4j in the future. We assume that you are already familiar with any general purpose programming language and have

some familiarity with Neo4j.

The two volume set CCIS 775 and 776 constitutes the refereed proceedings of the First International Conference on Computational Intelligence, Communications, and Business Analytics, CICBA 2017, held in Kolkata, India, in March 2017. The 90 revised full papers presented in the two volumes were carefully reviewed and selected from 276 submissions. The papers are organized in topical sections on data science and advanced data analytics; signal processing and communications; microelectronics, sensors, intelligent networks; computational forensics (privacy and security); computational intelligence in bio-computing; computational intelligence in mobile and quantum computing; intelligent data mining and data warehousing; computational intelligence.

NoSQL Starter is a great resource for someone starting with NoSQL and an indispensable guide for technology decision makers. It is assumed that you have a background in RDBMS modeling and SQL and have had exposure to at least one of the programming languages – Java or JavaScript. Friendly, practical tutorial with lots of hints and tips from several experienced Solr users and developers.

Learn advanced analytical techniques and leverage existing tool kits to make your analytic applications more powerful, precise, and efficient. This book provides the right combination of architecture, design, and implementation information to create analytical systems that go beyond the basics of classification, clustering, and recommendation. Pro Hadoop Data Analytics emphasizes best practices to ensure coherent, efficient development. A complete example system will be developed using standard third-party components that consist of the tool kits, libraries, visualization and reporting code, as well as support glue to provide a working and extensible end-to-end system. The book also highlights the importance of end-to-end, flexible, configurable, high-performance data pipeline systems with analytical components as well as appropriate visualization results. You'll discover the importance of mix-and-match or hybrid systems, using different analytical components in one application. This hybrid approach will be prominent in the examples. What You'll Learn Build big data analytic systems with the Hadoop ecosystem Use libraries, tool kits, and algorithms to make development easier and more effective Apply metrics to measure performance and efficiency of components and systems Connect to standard relational databases, noSQL data sources, and more Follow case studies with example components to create your own systems Who This Book Is For Software engineers, architects, and data scientists with an interest in the design and implementation of big data analytical systems using Hadoop, the Hadoop ecosystem, and other associated technologies.

La Data Science è una disciplina che unisce competenze scientifiche e tecnologiche di un determinato settore applicativo, con la capacità di elaborare risultati complessi, interpretando direttamente i dati raccolti. Il Machine Learning è uno degli indirizzi di successo dell'Intelligenza Artificiale. Apprendere vuol dire evolvere. Questi temi sono analizzati nel corso del libro, sottolineando l'efficacia delle strategie, il fascino e l'espressività delle idee emergenti.

Serving as a flagship driver towards advance research in the area of Big Data platforms and applications, this book provides a platform for the dissemination of advanced topics of theory, research efforts and analysis, and implementation oriented on methods, techniques and performance evaluation. In 23 chapters, several important formulations of the architecture design, optimization techniques, advanced analytics methods, biological, medical and social media applications are presented. These chapters discuss the research of members from the ICT COST Action IC1406 High-Performance Modelling and Simulation for Big Data Applications (cHiPSet). This volume is ideal as a reference for students, researchers and industry practitioners working in or interested in joining interdisciplinary works in the areas of intelligent decision systems using emergent distributed computing paradigms. It will also allow newcomers to grasp the key concerns and their potential solutions.

Quickly and productively develop complex Spring applications and microservices out of the box, with minimal concern over things like configurations. This revised book will show you how to fully leverage the Spring Boot 2 technology and how to apply it to create enterprise ready applications that just work. It will also cover what's been added to the new Spring Boot 2 release, including Spring Framework 5 features like WebFlux, Security, Actuator and the new way to expose Metrics through Micrometer framework, and more. This book is your authoritative hands-on practical guide for increasing your enterprise Java and cloud application productivity while decreasing development time. It's a no nonsense guide with case studies of increasing complexity throughout the book. The author, a senior solutions architect and Principal Technical instructor with Pivotal, the company behind the Spring Framework, shares his experience, insights and first-hand knowledge about how Spring Boot technology works and best practices. Pro Spring Boot 2 is an essential book for your Spring learning and reference library. What You Will Learn Configure and use Spring Boot Use non-functional requirements with Spring Boot Actuator Carry out web development with Spring Boot Persistence with JDBC, JPA and NoSQL Databases Messaging with JMS, RabbitMQ and WebSockets Test and deploy with Spring Boot A quick look at the Spring Cloud projects Microservices and deployment to the Cloud Extend Spring Boot by creating your own Spring Boot Starter and @Enable feature Who This Book Is For Experienced Spring and Java developers seeking increased productivity gains and decreased complexity and development time in their applications and software services.

If you are not already one of the 4 million plus users of the MySQL database, you will want to become one after reading this book. MySQL Administrator's Guide is the official guide to installing, administering and working with MySQL 5.0 and MySQL databases. Starting with the basics, you will learn to store, manage and retrieve data in a MySQL database and to eventually manipulate the code behind MySQL 5.0 to customize your database. With the clear examples and tips provided in this comprehensive learning guide, you can become a MySQL expert in no time.

With this practical book, architects, CTOs, and CIOs will learn a set of patterns for the practice of architecture, including analysis, documentation, and communication. Author Eben Hewitt shows you how to create holistic and thoughtful technology plans, communicate them clearly, lead people toward the vision, and become a great architect or Chief Architect. This book covers each key aspect of architecture comprehensively, including how to incorporate business architecture, information architecture, data architecture, application (software) architecture together to have the best chance for the system's success. Get a practical set of proven architecture practices focused on shipping great products using architecture Learn how architecture works effectively with development teams, management, and product management teams through the value chain Find updated special coverage on machine learning architecture Get usable templates to start incorporating into your teams immediately Incorporate business architecture, information architecture, data architecture, and application (software) architecture together This book constitutes the refereed proceedings of the 17th International Conference on Model Driven Engineering Languages and Systems, MODELS 2014, held in Valencia, Spain, in September/October 2014. The 41 full papers presented in this volume were carefully reviewed and selected from a total of 126 submissions. The scope of the conference series is broad, encompassing modeling languages, methods, tools, and applications considered from theoretical and practical angles and in academic and industrial settings. The papers report on the use of modeling in a wide range of cloud, mobile, and web computing, model transformation behavioral modeling, MDE: past, present, future, formal semantics, specification, and verification, models at runtime, feature and variability modeling, composition and adaptation, practices and experience, modeling for analysis, pragmatics, model extraction, manipulation and persistence, querying, and reasoning.

#### Getting Started with OrientDB

Implement DevOps with ease About This Book Leverage OneOps to achieve continuous application lifecycle management Switch between multiple cloud providers in order to leverage better pricing, technology, and scalability Build complex environments in a repeatable and predictable method, and deploy and scale on multiple clouds Who This Book Is For This book targets DevOps who want to use OneOps daily to deploy their applications, and Sysadmins who will be administering those applications. What You Will Learn Learn how to install OneOps Configure OneOps, including customizing your organizations, teams and clouds Work through practical deployment scenarios Understand OneOps architecture and individual components like Circuit and Display Build custom components and add unsupported clouds programmatically to OneOps Extend OneOps by calling the REST API In Detail Walmart's OneOps is an open source DevOps platform that is used for cloud and application lifecycle management. It can manage critical and complex application workload on any multi cloud-based infrastructure and revolutionizes the way administrators, developers, and engineers develop and launch new products. This practical book focuses on real-life cases and hands-on scenarios to develop, launch, and test your applications faster, so you can implement the DevOps process using OneOps. You will be exposed to the fundamental aspects of OneOps starting with installing, deploying, and configuring OneOps in a test environment, which will also come in handy later for development and debugging. You will also learn about design and architecture, and work through steps to perform enterprise level deployment. You will understand the initial setup of OneOps such as creating organization, teams, and access management. Finally, you will be taught how to configure, repair, scale, and extend applications across various cloud platforms. Style and approach A practical book packed with real life examples, practical scenarios and case studies on leveraging OneOps to implement DevOps for your organization.

Summary Neo4j in Action is a comprehensive guide to Neo4j, aimed at application developers and software architects. Using hands-on examples, you'll learn to model graph domains naturally with Neo4j graph structures. The book explores the full power of native Java APIs for graph data manipulation and querying. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Much of the data today is highly connected—from social networks to supply chains to software dependency management—and more connections are continually being uncovered. Neo4j is an ideal graph database tool for highly connected data. It is mature, production-ready, and unique in enabling developers to simply and efficiently model and query connected data. About the Book Neo4j in Action is a comprehensive guide to designing, implementing, and querying graph data using Neo4j. Using hands-on examples, you'll learn to model graph domains naturally with Neo4j graph structures. The book explores the full power of native Java APIs for graph data manipulation and querying. It also covers Cypher, Neo4j's graph query language. Along the way, you'll learn how to integrate Neo4j into your domain-driven app using Spring Data Neo4j, as well as how to use Neo4j in standalone server or embedded modes. Knowledge of Java basics is required. No prior experience with graph data or Neo4j is assumed. What's Inside Graph database patterns How to model data in social networks How to use Neo4j in your Java applications How to configure and set up Neo4j About the Authors Aleksa Vukotic is an architect specializing in graph data models. Nicki Watt, Dominic Fox, Tareq Abedrabbo, and Jonas Partner work at OpenCredo, a Neo Technology partner, and have been involved in many projects using Neo4j. Table of Contents PART 1 INTRODUCTION TO NEO4J A case for a Neo4j database Data modeling in Neo4j Starting development with Neo4j The power of traversals Indexing the data PART 2 APPLICATION DEVELOPMENT WITH NEO4J Cypher: Neo4j query language Transactions Traversals in depth Spring Data Neo4j PART 3 NEO4J IN PRODUCTION Neo4j: embedded versus server mode

[Copyright: 547f026804fbc460e5a7aaa7c03ae045](#)