Apache Solr 4 Cookbook

This book is written in a Cookbook style with short recipes showing developers how to effectively implement EIP without breaking everything in the process. It is concise and to the point, and it helps developers get their data flowing between different components without the need to read through page upon page of theory, while also enabling the reader to learn how to create exciting new projects. Camel Enterprise Integration Cookbook is intended for developers who have some familiarity with Apache Camel and who want a quick lookup reference to practical, proven tips on how to perform common tasks. Every recipe also includes a summary and reference pointers for more details that make it easy for you to get a deeper understanding of the Apache Camel capabilities that you will use day to day.

This book is for Elasticsearch users who want to extend their knowledge and develop new skills. Prior knowledge of the Query DSL and data indexing is expected.

This book is part of Packt's Cookbook series; each chapter looks at a different aspect of working with Apache Solr. The recipes deal with common problems of working with Solr by using easy-to-understand, real-life examples. The book is not in any way a complete Apache Solr reference and you should see it as a helping hand when things get rough on your journey with Apache Solr. Developers who are working with Apache Solr and would like to know how to combat common problems will find this book of great use. Knowledge of Apache Lucene would be a bonus but is not required.

Summary Redis in Action introduces Redis and walks you through examples that demonstrate how to use it effectively. You'll begin by getting Redis set up properly and then exploring the key-value model. Then, you'll dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets. Experienced developers will appreciate chapters on clustering and internal scripting to make Redis easier to use. About the Technology When you need near-real-time access to a fast-moving data stream, key-value stores like Redis are the way to go. Redis expands on the key-value pattern by accepting a wide variety of data types, including hashes, strings, lists, and other structures. It provides lightning-fast operations on in-memory datasets, and also makes it easy to persist to disk on the fly. Plus, it's free and open source. About this book Redis in Action introduces Redis and the key-value model. You'll quickly dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets and discover how to integrate with traditional RDBMS or other NoSQL stores. Experienced developers will appreciate the in-depth chapters on clustering and internal scripting. Written for developers familiar with database concepts. No prior exposure to NoSQL database concepts nor to Redis itself is required. Appropriate for systems administrators comfortable with programming. Purchase of the print book

includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside Redis from the ground up Preprocessing real-time data Managing in-memory datasets Pub/sub and configuration Persisting to disk About the Author Dr. Josiah L. Carlson is a seasoned database professional and an active contributor to the Redis community. Table of Contents PART 1 GETTING STARTED Getting to know Redis Anatomy of a Redis web application PART 2 CORE CONCEPTS Commands in Redis Keeping data safe and ensuring performance Using Redis for application support Application components in Redis Search-based applications Building a simple social network PART 3 NEXT STEPS Reducing memory use Scaling Redis Scripting Redis with Lua

ElasticSearch is an open source search server built on Apache Lucene. It was built to provide a scalable search solution with built-in support for near real-time search and multi-tenancy.Jumping into the world of ElasticSearch by setting up your own custom cluster, this book will show you how to create a fast, scalable, and flexible search solution. By learning the ins-and-outs of data indexing and analysis, "ElasticSearch Server" will start you on your journey to mastering the powerful capabilities of ElasticSearch. With practical chapters covering how to search data, extend your search, and go deep into cluster administration and search analysis, this book is perfect for those new and experienced with search servers.In "ElasticSearch Server" you will learn how to revolutionize your website or application with faster, more accurate, and flexible search functionality. Starting with chapters on setting up your own ElasticSearch cluster and searching and extending your search parameters you will quickly be able to create a fast, scalable, and completely custom search solution.Building on your knowledge further you will learn about ElasticSearch's query API and become confident using powerful filtering and faceting capabilities. You will develop practical knowledge on how to make use of ElasticSearch's near real-time capabilities and support for multi-tenancy.Your journey then concludes with chapters that help you monitor and tune your ElasticSearch cluster as well as advanced topics such as shard allocation, gateway configuration, and the discovery module.

JSON is becoming the backbone for meaningful data interchange over the internet. This format is now supported by an entire ecosystem of standards, tools, and technologies for building truly elegant, useful, and efficient applications. With this hands-on guide, author and architect Tom Marrs shows you how to build enterprise-class applications and services by leveraging JSON tooling and message/document design. JSON at Work provides application architects and developers with guidelines, best practices, and use cases, along with lots of real-world examples and code samples. You'll start with a comprehensive JSON overview, explore the JSON ecosystem, and then dive into JSON's use in the enterprise. Get acquainted with JSON basics and learn how to model JSON data Learn how to use JSON with Node.js, Ruby on Rails, and Java Structure JSON documents with JSON Schema to design and test APIs Search the contents of

JSON documents with JSON Search tools Convert JSON documents to other data formats with JSON Transform tools Compare JSON-based hypermedia formats, including HAL and jsonapi Leverage MongoDB to store and access JSON documents Use Apache Kafka to exchange JSON-based messages between services

This book is for intermediate Solr Developers who are willing to learn and implement Pro-level practices, techniques, and solutions. This edition will specifically appeal to developers who wish to quickly get to grips with the changes and new features of Apache Solr 5.

Building distributed applications is difficult enough without having to coordinate the actions that make them work. This practical guide shows how Apache ZooKeeper helps you manage distributed systems, so you can focus mainly on application logic. Even with ZooKeeper, implementing coordination tasks is not trivial, but this book provides good practices to give you a head start, and points out caveats that developers and administrators alike need to watch for along the way. In three separate sections, ZooKeeper contributors Flavio Junqueira and Benjamin Reed introduce the principles of distributed systems, provide ZooKeeper programming techniques, and include the information you need to administer this service. Learn how ZooKeeper solves common coordination tasks Explore the ZooKeeper API’s Java and C implementations and how they differ Use methods to track and react to ZooKeeper state changes Handle failures of the network, application processes, and ZooKeeper itself Learn about ZooKeeper’s trickier aspects dealing with concurrency, ordering, and configuration Use the Curator high-level interface for connection management Become familiar with ZooKeeper internals and administration tools

Recipes to help you overcome your data science hurdles using Java About This Book This book provides modern recipes in small steps to help an apprentice cook become a master chef in data science Use these recipes to obtain, clean, analyze, and learn from your data Learn how to get your data science applications to production and enterprise environments effortlessly Who This Book Is For This book is for Java developers who are familiar with the fundamentals of data science and want to improve their skills to become a pro. What You Will Learn Find out how to clean and make datasets ready so you can acquire actual insights by removing noise and outliers Develop the skills to use modern machine learning techniques to retrieve information and transform data to knowledge. retrieve information from large amount of data in text format. Familiarize yourself with cutting-edge techniques to store and search large volumes of data and retrieve information from large amounts of data in text format Develop basic skills to apply big data and deep learning technologies on large volumes of data Evolve your data visualization skills and gain valuable insights from your data Get to know a step-by-step formula to develop an industry-standard, large-scale, real-life data product Gain the skills to visualize data and interact with users through data insights In Detail If you are looking to build data science models that are good for production, Java has come to the rescue. With the aid of strong libraries such as MLlib, Weka, DL4j, and more, you can efficiently perform all the data science tasks you need to. This unique book provides modern recipes to solve your common and not-so-common data science-related problems. We start with recipes to help you obtain, clean, index, and search data. Then you will learn a variety of techniques to analyze, learn from, and retrieve information from data. You will also understand how to handle big

data, learn deeply from data, and visualize data. Finally, you will work through unique recipes that solve your problems while taking data science to production, writing distributed data science applications, and much more—things that will come in handy at work. Style and approach This book contains short yet very effective recipes to solve most common problems. Some recipes cater to very specific, rare pain points. The recipes cover different data sets and work very closely to real production environments

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

Whether you're deploying applications on-premise or in the cloud, this cookbook is for developers, operators, and IT professionals who need practical solutions for using Docker. The recipes in this book will help developers go from zero knowledge to distributed applications packaged and deployed within a couple of chapters. IT professionals will be able to use this cookbook to solve everyday problems, as well as create, run, share, and deploy Docker images quickly. Operators will learn and understand what developers are excited about and start to adopt the tools that will change the way they work.--

Summary CMIS and Apache Chemistry in Action is a comprehensive guide to the CMIS standard and related ECM concepts, written by the authors of the standard. In it, you'll tackle hands-on examples for building applications on CMIS repositories from both the client and the server sides. You'll learn how to create new content-centric applications that install and run in any CMIS-compliant repository. About The Technology Content Management Interoperability Services (CMIS) is an OASIS standard for accessing content management systems. It specifies a vendor-and language-neutral way to interact with any compliant content repository. Apache Chemistry provides complete reference implementations of the CMIS standard with robust APIs for developers writing tools, applications, and servers. About This Book CMIS and Apache Chemistry in Action is a comprehensive guide to the CMIS standard and related ECM concepts. In it, you'll find clear teaching and instantly useful examples for building content-centric client and server-side applications that run against any CMIS-compliant repository. In fact, using the CMIS Workbench and the InMemory Repository from Apache Chemistry, you'll have running code talking to a real CMIS server by the end of chapter 1. This book requires some familiarity with content management systems and a standard programming language like Java or C#. No exposure to CMIS or Apache Chemistry is assumed. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside The only CMIS book endorsed by OASIS Complete coverage of the CMIS 1.0 and 1.1 specifications Cookbook-style tutorials and real-world examples About the Authors Florian Müller, Jay Brown, and Jeff Potts are among the original authors, contributors, and leaders of Apache Chemistry and the OASIS CMIS specification. They continue to shape CMIS implementations at Alfresco, IBM, and SAP. Table of Contents PART 1 UNDERSTANDING CMIS Introducing CMIS Exploring the CMIS domain model Creating, updating, and deleting objects with CMIS CMIS metadata: types and properties Query PART 2 HANDS-ON CMIS CLIENT DEVELOPMENT Meet your new project: The Blend The Blend: read and query functionality The Blend: create, update, and delete functionality Using other client libraries Building mobile apps with CMIS PART 3 ADVANCED TOPICS CMIS bindings Security and control Performance Building a CMIS server

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.

Build an enterprise search engine using Apache Solr: index and search documents; ingest data from varied sources; apply various text processing techniques; utilize different search capabilities; and customize Solr to retrieve the desired results. Apache Solr: A Practical Approach to Enterprise Search explains each essential concept-backed by practical and industry examples--to help you attain expert-level knowledge. The book, which assumes a basic knowledge of Java, starts with an introduction to Solr, followed by steps to setting it up, indexing your first set of documents, and searching them. It then introduces you to information retrieval and its implementation in Apache Solr; this will help you understand your search problem, decide the approach to build an effective solution, and use various metrics to evaluate the results. The book next covers the schema design and techniques to build a text analysis chain for cleansing, normalizing and enriching your documents and addressing different types of search queries. It describes various popular matching techniques which are generally applied to improve the precision and recall of searches. You will learn the end-to-end process of data ingestion from varied sources, metadata extraction, pre-processing and transformation of content, various search components, query parsers and other advanced search capabilities. After covering out-of-the-box features, Solr expert Dikshant Shahi dives into ways you can customize Solr for your business and its specific requirements, along with ways to plug in your own components. Most important, you will learn about implementations for Solr scoring, factors affecting the document score, and tuning the score for the application at hand. The book explains why textual scoring is not sufficient for practical ranking of documents and ways to integrate real-world factors for contributing to the document ranking. You'll see how to influence user experience by providing suggestions and recommendations. You'll also see integration of Solr with important related technologies such as OpenNLP and Tika. Additionally, you will learn about scaling Solr using SolrCloud. This book concludes with coverage of semantic search capabilities, which is crucial for taking the search experience to the next level. By the end of Apache Solr, you will be proficient in designing and developing your search engine.

This book is for developers who want to learn how to get the most out of Solr in their applications, whether you are new to the field, have used Solr but don't know everything, or simply want a good reference. It would be helpful to have some familiarity with basic programming concepts, but no prior experience is required.

Summary A developer-focused guide to writing applications using Spring Boot. You'll learn how to bypass the tedious configuration steps so

that you can concentrate on your application's behavior. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The Spring Framework simplifies enterprise Java development, but it does require lots of tedious configuration work. Spring Boot radically streamlines spinning up a Spring application. You get automatic configuration and a model with established conventions for build-time and runtime dependencies. You also get a handy command-line interface you can use to write scripts in Groovy. Developers who use Spring Boot often say that they can't imagine going back to hand configuring their applications. About the Book Spring Boot in Action is a developer-focused guide to writing applications using Spring Boot. In it, you'll learn how to bypass configuration steps so you can focus on your application's behavior. Spring expert Craig Walls uses interesting and practical examples to teach you both how to use the default settings effectively and how to override and customize Spring Boot for your unique environment. Along the way, you'll pick up insights from Craig's years of Spring development experience. What's Inside Develop Spring apps more efficiently Minimal to no configuration Runtime metrics with the Actuator Covers Spring Boot 1.3 About the Reader Written for readers familiar with the Spring Framework. About the Author Craig Walls is a software developer, author of the popular book Spring in Action, Fourth Edition, and a frequent speaker at conferences. Table of Contents Bootstarting Spring Developing your first Spring Boot application Customizing configuration Testing with Spring Boot Getting Groovy with the Spring Boot CLI Applying Grails in Spring Boot Taking a peek inside with the Actuator Deploying Spring Boot applications APPENDIXES Spring Boot developer tools Spring Boot starters Configuration properties Spring Boot dependencies

Lucene 4 Cookbook is a practical guide that shows you how to build a scalable search engine for your application, from an internal documentation search to a wide-scale web implementation with millions of records. Starting with helping you to successfully install Apache Lucene, it will guide you through creating your first search application. Furthermore, the book walks you through analyzing your text and indexing your data to leverage the performance of your search application. As you progress through the chapters, you will learn to effectively search your indexes and successfully employ real-time searching. The chapters start off with simple concepts and build up to complex solutions that should help you on your way to becoming a search engine expert.

If you are a developer who implements ElasticSearch in your web applications and want to sharpen your understanding of the core elements and applications, this is the book for you. It is assumed that you’ve got working knowledge of JSON and, if you want to extend ElasticSearch, of Java and related technologies.

Whether you need full-text search or real-time analytics of structured data—or both—the Elasticsearch distributed search engine is an ideal way to put your data to work. This practical guide not only shows you how to search, analyze, and explore data with Elasticsearch, but also helps you deal with the complexities of human language, geolocation, and relationships. If you’re a newcomer to both search and distributed systems, you’ll quickly learn how to integrate Elasticsearch into your application. More experienced users will pick up lots of advanced techniques. Throughout the book, you’ll follow a problem-based approach to learn why, when, and how to use Elasticsearch features. Understand how Elasticsearch interprets data in your documents Index and query your data to take advantage of search concepts such as relevance and word proximity Handle human language through the effective use of analyzers and queries Summarize and group data to show overall trends, with aggregations and analytics Use geo-points and geo-shapes—Elasticsearch’s approaches to geolocation Model your data to take advantage of Elasticsearch’s horizontal scalability Learn how to configure and monitor your cluster in production

This book is in Packt's Cookbook series. A Packt Cookbook contains recipes for solutions to the most important problems you face when

working with a topic. Inside the Cookbook you will find: A straightforward and easy to follow format, A selection of the most important tasks and problems, Carefully organized instructions for solving the problem efficiently, Clear explanations of what you did, Details for applying the solution to other situations. This is the ideal book for people who have already written a first application with the Play Framework or have just finished reading through the documenation. In other words - anyone who is ready to get to grips with Play! Having a basic knowledge of Java is good, as well as well as some web developer skills – HTML and JavaScript

Over 50 hands-on recipes to efficiently administer, maintain, and use your Apache Kafka installationAbout This Book- Quickly configure and manage your Kafka cluster- Learn how to use the Apache Kafka cluster and connect it with tools for big data processing- A practical guide to monitor your Apache Kafka installationWho This Book Is ForIf you are a programmer or big data engineer using or planning to use Apache Kafka, then this book is for you. This book has several recipes which will teach you how to effectively use Apache Kafka. You need to have some basic knowledge of Java. If you don't know big data tools, this would be your stepping stone for learning how to consume the data in these kind of systems.What You Will Learn- Learn how to configure Kafka brokers for better efficiency- Explore how to configure producers and consumers for optimal performance- Set up tools for maintaining and operating Apache Kafka- Create producers and consumers for Apache Kafka in Java- Understand how Apache Kafka can be used by several third party system for big data processing, such as Apache Storm, Apache Spark, Hadoop, and more- Monitor Apache Kafka using tools like graphite and GangliaIn DetailThis book will give you details about how to manage and administer your Apache Kafka Cluster.We will cover topics like how to configure your broker, producer, and consumer for maximum efficiency for your situation. Also, you will learn how to maintain and administer your cluster for fault tolerance. We will also explore tools provided with Apache Kafka to do regular maintenance operations. We shall also look at how to easily integrate Apache Kafka with big data tools like Hadoop, Apache Spark, Apache Storm, and Elasticsearch.Style and approachEasy-to-follow, step-by-step recipes explaining from start to finish how to accomplish real-world tasks.

Over 60 recipes to help you speed up the development of your Java web applications using the Spring Roo development tool.

What could you do with data if scalability wasn't a problem? With this hands-on guide, you'll learn how Apache Cassandra handles hundreds of terabytes of data while remaining highly available across multiple data centers -- capabilities that have attracted Facebook, Twitter, and other data-intensive companies. Cassandra: The Definitive Guide provides the technical details and practical examples you need to assess this database management system and put it to work in a production environment. Author Eben Hewitt demonstrates the advantages of Cassandra's nonrelational design, and pays special attention to data modeling. If you're a developer, DBA, application architect, or manager looking to solve a database scaling issue or future-proof your application, this guide shows you how to harness Cassandra's speed and flexibility. Understand the tenets of Cassandra's column-oriented structure Learn how to write, update, and read Cassandra data Discover how to add or remove nodes from the cluster as your application requires Examine a working application that translates from a relational model to Cassandra's data model Use examples for writing clients in Java, Python, and C# Use the JMX interface to monitor a cluster's usage, memory patterns, and more Tune memory settings, data storage, and caching for better performance

Search, analyze, and manage data effectively with Elasticsearch 7 Key Features Extend Elasticsearch functionalities and learn how to deploy on Elastic Cloud Deploy and manage simple Elasticsearch nodes as well as complex cluster topologies Explore the capabilities of Elasticsearch 7 with easy-to-follow recipes Book Description Elasticsearch is a Lucene-based distributed search server that allows users to index and search unstructured content with petabytes of data. With this book, you'll be guided through comprehensive recipes on what's new

in Elasticsearch 7, and see how to create and run complex queries and analytics. Packed with recipes on performing index mapping, aggregation, and scripting using Elasticsearch, this fourth edition of Elasticsearch Cookbook will get you acquainted with numerous solutions and quick techniques for performing both every day and uncommon tasks such as deploying Elasticsearch nodes, integrating other tools to Elasticsearch, and creating different visualizations. You will install Kibana to monitor a cluster and also extend it using a variety of plugins. Finally, you will integrate your Java, Scala, Python, and big data applications such as Apache Spark and Pig with Elasticsearch, and create efficient data applications powered by enhanced functionalities and custom plugins. By the end of this book, you will have gained in-depth knowledge of implementing Elasticsearch architecture, and you'll be able to manage, search, and store data efficiently and effectively using Elasticsearch. What you will learn Create an efficient architecture with Elasticsearch Optimize search results by executing analytics aggregations Build complex queries by managing indices and documents Monitor the performance of your cluster and nodes Design advanced mapping to take full control of index steps Integrate Elasticsearch in Java, Scala, Python, and big data applications Install Kibana to monitor clusters and extend it for plugins Who this book is for If you’re a software engineer, big data infrastructure engineer, or Elasticsearch developer, you'll find this book useful. This Elasticsearch book will also help data professionals working in the e-commerce and FMCG industry who use Elastic for metrics evaluation and search analytics to get deeper insights for better business decisions. Prior experience with Elasticsearch will help you get the most out of this book.

Summary MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology This document-oriented database was built for high availability, supports rich, dynamic schemas, and lets you easily distribute data across multiple servers. MongoDB 3.0 is flexible, scalable, and very fast, even with big data loads. About the Book MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Lots of examples will help you develop confidence in the crucial area of data modeling. You'll also love the deep explanations of each feature, including replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engineer, Doug Garrett is one of the winners of the MongoDB Innovation Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihalis Tsoukalos. Table of Contents PART 1 GETTING STARTED A database for the modern web

MongoDB through the JavaScript shell Writing programs using MongoDB PART 2 APPLICATION DEVELOPMENT IN MONGODB Document-oriented data Constructing queries Aggregation Updates, atomic operations, and deletes PART 3 MONGODB MASTERY Indexing and query optimization Text search WiredTiger and pluggable storage Replication Scaling your system with sharding Deployment and administration

Learn everything you need to build highly scalable, robust web applications using Angular release 4 About This Book Apply best practices and design patterns to achieve higher scalability in your Angular applications Understand the latest features of Angular and create your own components Get acquainted with powerful, advanced techniques in Angular to build professional web applications Who This Book Is For This book is for JavaScript developers with some prior exposure to Angular, at least through basic examples. We assume that you've got working knowledge of HTML, CSS, and JavaScript. What You Will Learn Implement asynchronous programming using Angular Beautify your application with the UI components built to the material design specification Secure your web application from unauthorized users Create complex forms, taking full advantage of 2-way data binding Test your Angular applications using the Jasmine and Protractor frameworks for better efficiency Learn how to integrate Angular with Bootstrap to create compelling web applications Use Angular built-in classes to apply animation in your app In Detail Got some experience of Angular under your belt? Want to learn everything about using advanced features for developing websites? This book is everything you need for the deep understanding of Angular that will set you apart from the developer crowd. Angular has introduced a new way to build applications. Creating complex and rich web applications, with a lighter resource footprint, has never been easier or faster. Angular is now at release 4, with significant changes through previous versions. This book has been written and tested for Angular release 4. Angular is a mature technology, and you'll likely have applications built with earlier versions. This book starts by showing you best practices and approaches to migrating your existing Angular applications so that you can be immediately up-to-date. You will take an in-depth look at components and see how to control the user journey in your applications by implementing routing and navigation. You will learn how to work with asynchronous programming by using Observables. To easily build applications that look great, you will learn all about template syntax and how to beautify applications with Material Design. Mastering forms and data binding will further speed up your application development time. Learning about managing services and animations will help you to progressively enhance your applications. Next you'll use native directives to integrate Bootstrap with Angular. You will see the best ways to test your application with the leading options such as Jasmine and Protractor. At the end of the book, you'll learn how to apply design patterns in Angular, and see the benefits they will bring to your development. Style and approach This book provides comprehensive coverage of all aspects of development with Angular. You will learn

about all the most powerful Angular concepts, with examples and best practices. This book is everything you need for the deep understanding of Angular that will set you apart from the developer crowd.

From lambda expressions and JavaFX 8 to new support for network programming and mobile development, Java 8 brings a wealth of changes. This cookbook helps you get up to speed right away with hundreds of hands-on recipes across a broad range of Java topics. You'll learn useful techniques for everything from debugging and data structures to GUI development and functional programming. Each recipe includes self-contained code solutions that you can freely use, along with a discussion of how and why they work. If you are familiar with Java basics, this cookbook will bolster your knowledge of the language in general and Java 8's main APIs in particular. Recipes include: Methods for compiling, running, and debugging Manipulating, comparing, and rearranging text Regular expressions for string- and pattern-matching Handling numbers, dates, and times Structuring data with collections, arrays, and other types Object-oriented and functional programming techniques Directory and filesystem operations Working with graphics, audio, and video GUI development, including JavaFX and handlers Network programming on both client and server Database access, using JPA, Hibernate, and JDBC Processing JSON and XML for data storage Multithreading and concurrency

Apache Mahout Cookbook uses over 35 recipes packed with illustrations and real-world examples to help beginners as well as advanced programmers get acquainted with the features of Mahout."Apache Mahout Cookbook" is great for developers who want to have a fresh and fast introduction to Mahout coding. No previous knowledge of Mahout is required, and even skilled developers or system administrators will benefit from the various recipes presented

Apache Solr 4 CookbookPackt Publishing Ltd

Enterprise and web applications require full-featured, "Google-quality" searchcapabilities, but such features are notoriously difficult to implement and maintain.Hibernate Search builds on the Lucene feature set and offers an easyto-implement interface that integrates seamlessly with Hibernate-the leadingdata persistence solution for Java applications. Hibernate Search in Action introduces both the principles of enterprise searchand the implementation details a Java developer will need to use HibernateSearch effectively. This book blends the insights of the Hibernate Search leaddeveloper with the practical techniques required to index and manipulate data,assemble and execute search queries, and create smart filters for better searchresults. Along the way, the reader masters performance-boosting concepts likeusing Hibernate Search in a clustered environment and integrating with thefeatures already in your applications. This book assumes you're a competent Java developer with some experienceusing Hibernate and Lucene. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.

This book will provide clear guidance on how to work through the most valuable design patterns effectively in Angular. You will explore some of the best ways to work with Angular to meet the performance required in the web development world. You will also learn the best practices to improve your productivity and the code base of your application.
Ansible is a simple, but powerful, server and configuration management tool. Learn to use Ansible effectively, whether you manage one server--or thousands.
"Apache Solr 4 Cookbook" is written in a helpful, practical style with numerous hands-on recipes to help you master Apache Solr to get more precise search results and analysis, higher performance, and reliability. This book is for developers who wish to learn how to master Apache Solr 4. This book will specifically appeal to developers who wish to quickly get to grips with the changes and new features of Apache Solr 4. This book is also handy as a practical guide to solving common problems and issues when using Apache Solr.
Over 100 practical recipes to make Apache Solr faster, more reliable and return better results.
Implement natural language processing applications with Python using a problem-solution approach. This book has numerous coding exercises that will help you to quickly deploy natural language processing techniques, such as text classification, parts of speech identification, topic modeling, text summarization, text generation, entity extraction, and sentiment analysis. Natural Language Processing Recipes starts by offering solutions for cleaning and preprocessing text data and ways to analyze it with advanced algorithms. You'll see practical applications of the semantic as well as syntactic analysis of text, as well as complex natural language processing approaches that involve text normalization, advanced preprocessing, POS tagging, and sentiment analysis. You will also learn various applications of machine learning and deep learning in natural language processing. By using the recipes in this book, you will have a toolbox of solutions to apply to your own projects in the real world, making your development time quicker and more efficient. What You Will Learn Apply NLP techniques using Python libraries such as NLTK, TextBlob, spaCy, Stanford CoreNLP, and many more Implement the concepts of information retrieval, text summarization, sentiment analysis, and other advanced natural language processing techniques. Identify machine learning and deep learning techniques for natural language processing and natural language generation problems Who This Book Is ForData scientists who want to refresh and learn various concepts of natural language processing through coding exercises.
Individual self-contained code recipes. Solve specific problems using individual recipes, or work through the book to develop your capabilities. If you are a big data enthusiast and striving to use Hadoop to solve your problems, this book is for you. Aimed at Java programmers with some knowledge of Hadoop MapReduce, this is also a comprehensive reference for developers and system admins who want to get up to speed using Hadoop.
A fast paced guide that will help you learn about Apache Hadoop 3 and its ecosystem Key Features Set up, configure and get started with Hadoop to get useful insights from large data sets Work with the different components of Hadoop such as MapReduce, HDFS and YARN Learn about the new features introduced in Hadoop 3 Book Description Apache Hadoop is a widely used distributed data platform. It enables large datasets to be efficiently processed instead of using one large computer to store and process the data. This book will get you started with the Hadoop ecosystem, and introduce you to the main technical topics, including MapReduce, YARN, and HDFS. The book begins with an overview of big data and Apache Hadoop. Then, you will set up a pseudo Hadoop development environment and a multi-node enterprise Hadoop cluster. You will see how the parallel programming paradigm, such as MapReduce, can solve many complex data processing

problems. The book also covers the important aspects of the big data software development lifecycle, including quality assurance and control, performance, administration, and monitoring. You will then learn about the Hadoop ecosystem, and tools such as Kafka, Sqoop, Flume, Pig, Hive, and HBase. Finally, you will look at advanced topics, including real time streaming using Apache Storm, and data analytics using Apache Spark. By the end of the book, you will be well versed with different configurations of the Hadoop 3 cluster. What you will learn Store and analyze data at scale using HDFS, MapReduce and YARN Install and configure Hadoop 3 in different modes Use Yarn effectively to run different applications on Hadoop based platform Understand and monitor how Hadoop cluster is managed Consume streaming data using Storm, and then analyze it using Spark Explore Apache Hadoop ecosystem components, such as Flume, Sqoop, HBase, Hive, and Kafka Who this book is for Aspiring Big Data professionals who want to learn the essentials of Hadoop 3 will find this book to be useful. Existing Hadoop users who want to get up to speed with the new features introduced in Hadoop 3 will also benefit from this book. Having knowledge of Java programming will be an added advantage.

This book is aimed at developers, designers, and architects who would like to build big data enterprise search solutions for their customers or organizations. No prior knowledge of Apache Hadoop and Apache Solr/Lucene technologies is required.

A comprehensive guide to mastering the most advanced Hadoop 3 concepts Key Features Get to grips with the newly introduced features and capabilities of Hadoop 3 Crunch and process data using MapReduce, YARN, and a host of tools within the Hadoop ecosystem Sharpen your Hadoop skills with real-world case studies and code Book Description Apache Hadoop is one of the most popular big data solutions for distributed storage and for processing large chunks of data. With Hadoop 3, Apache promises to provide a high-performance, more fault-tolerant, and highly efficient big data processing platform, with a focus on improved scalability and increased efficiency. With this guide, you'll understand advanced concepts of the Hadoop ecosystem tool. You'll learn how Hadoop works internally, study advanced concepts of different ecosystem tools, discover solutions to real-world use cases, and understand how to secure your cluster. It will then walk you through HDFS, YARN, MapReduce, and Hadoop 3 concepts. You'll be able to address common challenges like using Kafka efficiently, designing low latency, reliable message delivery Kafka systems, and handling high data volumes. As you advance, you'll discover how to address major challenges when building an enterprise-grade messaging system, and how to use different stream processing systems along with Kafka to fulfil your enterprise goals. By the end of this book, you'll have a complete understanding of how components in the Hadoop ecosystem are effectively integrated to implement a fast and reliable data pipeline, and you'll be equipped to tackle a range of real-world problems in data pipelines. What you will learn Gain an in-depth understanding of distributed computing using Hadoop 3 Develop enterprise-grade applications using Apache Spark, Flink, and more Build scalable and high-performance Hadoop data pipelines with security, monitoring, and data governance Explore batch data processing patterns and how to model data in Hadoop Master best practices for enterprises using, or planning to use, Hadoop 3 as a data platform Understand security aspects of Hadoop, including authorization and authentication Who this book is for If you want to become a big data professional by mastering the advanced concepts of Hadoop, this book is for you. You'll also find this book useful if you're a Hadoop professional looking to strengthen your knowledge of the Hadoop ecosystem. Fundamental knowledge of the Java programming language and basics of Hadoop is necessary to get started with this book.

Summary Solr in Action is a comprehensive guide to implementing scalable search using Apache Solr. This clearly written book walks you through well-documented examples ranging from basic keyword searching to scaling a system for billions of documents and queries. It will give you a deep understanding of how to implement core Solr capabilities. About the Book Whether you're handling big (or small) data,

managing documents, or building a website, it is important to be able to quickly search through your content and discover meaning in it. Apache Solr is your tool: a ready-to-deploy, Lucene-based, open source, full-text search engine. Solr can scale across many servers to enable real-time queries and data analytics across billions of documents. Solr in Action teaches you to implement scalable search using Apache Solr. This easy-to-read guide balances conceptual discussions with practical examples to show you how to implement all of Solr's core capabilities. You'll master topics like text analysis, faceted search, hit highlighting, result grouping, query suggestions, multilingual search, advanced geospatial and data operations, and relevancy tuning. This book assumes basic knowledge of Java and standard database technology. No prior knowledge of Solr or Lucene is required. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside How to scale Solr for big data Rich real-world examples Solr as a NoSQL data store Advanced multilingual, data, and relevancy tricks Coverage of versions through Solr 4.7 About the Authors Trey Grainger is a director of engineering at CareerBuilder. Timothy Potter is a senior member of the engineering team at LucidWorks. The authors work on the scalability and reliability of Solr, as well as on recommendation engine and big data analytics technologies. Table of Contents PART 1 MEET SOLR Introduction to Solr Getting to know Solr Key Solr concepts Configuring Solr Indexing Text analysis PART 2 CORE SOLR CAPABILITIES Performing queries and handling results Faceted search Hit highlighting Query suggestions Result grouping/field collapsing Taking Solr to production PART 3 TAKING SOLR TO THE NEXT LEVEL SolrCloud Multilingual search Complex query operations Mastering relevancy