

Elasticsearch In Action

Elasticsearch in Action *Elasticsearch in Action* **Elasticsearch: The Definitive Guide** *Relevant Search* **Mastering Elasticsearch - Second Edition** *Elasticsearch Server* **Solr in Action** **Learning Elastic Stack 7.0** *Advanced Elasticsearch 7.0* **Lucene in Action** **Spark in Action** *The Logstash Book* **Elasticsearch Essentials** **Learning Elasticsearch** **Monitoring Elasticsearch** *Elasticsearch 7.0 Cookbook* *Akka in Action* **Learning ELK Stack** *ElasticSearch Cookbook - Second Edition* **Docker in Action** *Elasticsearch in Action, Second Edition* **MongoDB in Action** **Elasticsearch 7 Quick Start Guide** *Learning Kibana 7* *Storm Applied* *Architecting and Operating OpenShift Clusters* **Restlet in Action** *Mastering Elastic Stack* *I Heart Logs* **Microservices in Action** **Learning Elasticsearch 7.x** *Apache Pulsar in Action* *Neo4j in Action* *Mastering Kibana 6.x* **Getting Started with Kubernetes** *Just Spring* **Spring Security in Action** *Mastering Search Analytics* **Sencha Touch in Action** **Deep Learning for Search**

As recognized, adventure as skillfully as experience very nearly lesson, amusement, as with ease as deal can be gotten by just checking out a book **Elasticsearch In Action** with it is not directly done, you could agree to even more as regards this life, with reference to the world.

We pay for you this proper as with ease as simple artifice to acquire those all. We pay for Elasticsearch In Action and numerous book collections from fictions to scientific research in any way. in the midst of them is this Elasticsearch In Action that can be your partner.

Learning Elastic Stack 7.0 Mar 28 2022 A beginner's guide to storing, managing, and analyzing data with the updated features of Elastic 7.0 Key Features Gain access to new features and updates introduced in Elastic Stack 7.0 Grasp the fundamentals of Elastic Stack including Elasticsearch, Logstash, and Kibana Explore useful tips for using Elastic Cloud and deploying Elastic Stack in production environments Book Description The Elastic Stack is a powerful combination of tools for techniques such as distributed search, analytics, logging, and visualization of data. Elastic Stack 7.0 encompasses new features and capabilities that will enable you to find unique insights into analytics using these techniques. This book will give you a fundamental understanding of what the stack is all about, and help you use it efficiently to build powerful real-time data processing applications. The first few sections of the book will help you understand how to set up the stack by installing tools, and exploring their basic configurations. You'll then get up to speed with using Elasticsearch for distributed searching and analytics, Logstash for logging, and Kibana for data visualization. As you work through the book, you will discover the technique of creating custom plugins using Kibana and Beats. This is followed by coverage of the Elastic X-Pack, a useful extension for effective security and monitoring. You'll also find helpful tips on how to use Elastic Cloud and deploy Elastic Stack in production environments. By the end of this book, you'll be well versed with the fundamental Elastic Stack functionalities and the role of each component in the stack to solve different data processing problems. What you will learn Install and configure an Elasticsearch architecture Solve the full-text search problem with Elasticsearch Discover powerful analytics capabilities through aggregations using Elasticsearch Build a data pipeline to transfer data from a variety of sources into Elasticsearch for analysis Create interactive dashboards for effective storytelling with your data using Kibana Learn how to secure, monitor and use Elastic Stack's alerting and reporting capabilities Take applications to an on-premise or cloud-based production environment with Elastic Stack Who this book is for This book is for entry-level data professionals, software engineers, e-commerce developers, and full-stack developers who want to learn about Elastic Stack and how the real-time processing and search engine works for business analytics and enterprise search applications. Previous experience with Elastic Stack is not required, however knowledge of data warehousing and database concepts will be helpful.

Elasticsearch: The Definitive Guide Sep 02 2022 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

Spring Security in Action Sep 29 2019 Spring Security in Action shows you how to prevent cross-site scripting and request forgery attacks before they do damage. You'll start with the basics, simulating password upgrades and adding multiple types of authorization. As your skills grow, you'll adapt Spring Security to new architectures and create advanced OAuth2 configurations. By the time you're done, you'll have a customized Spring Security configuration that protects against threats both common and extraordinary. Summary While creating secure applications is critically important, it can also be tedious and time-consuming to stitch together the required collection of tools. For Java developers, the powerful Spring Security framework makes it easy for you to bake security into your software from the very beginning. Filled with code samples and practical examples, Spring Security in Action teaches you how to secure your apps from the most common threats, ranging from injection attacks to lackluster monitoring. In it, you'll learn how to manage system users, configure secure endpoints, and use OAuth2 and OpenID Connect for authentication and authorization. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Security is non-negotiable. You rely on Spring applications to transmit data, verify credentials, and prevent attacks. Adopting "secure by design" principles will protect your network from data theft and unauthorized intrusions. About the book Spring Security in Action shows you how to prevent cross-site scripting and request forgery attacks before they do damage. You'll start with the basics, simulating password upgrades and adding multiple types of authorization. As your skills grow, you'll adapt Spring Security to new architectures and create advanced OAuth2 configurations. By the time you're done, you'll have a customized Spring Security configuration that protects against threats both common and extraordinary. What's inside Encoding passwords and authenticating users Securing endpoints Automating security testing Setting up a standalone authorization server About the reader For experienced Java and Spring developers. About the author Laurentiu Spilca is a dedicated development lead and trainer at Endava, with over ten years of Java experience. Table of Contents PART 1 - FIRST STEPS 1 Security Today 2 Hello Spring Security PART 2 - IMPLEMENTATION 3 Managing users 4 Dealing with passwords 5 Implementing authentication 6 Hands-on: A small secured web application 7 Configuring authorization: Restricting access 8 Configuring authorization: Applying restrictions 9 Implementing filters 10 Applying CSRF protection and CORS 11 Hands-on: A separation of responsibilities 12 How does OAuth 2 work? 13 OAuth 2: Implementing the authorization server 14 OAuth 2: Implementing the resource server

15 OAuth 2: Using JWT and cryptographic signatures 16 Global method security: Pre- and postauthorizations 17 Global method security: Pre- and postfiltering 18 Hands-on: An OAuth 2 application 19 Spring Security for reactive apps 20 Spring Security testing

Mastering Elasticsearch - Second Edition Jun 30 2022 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.

Learning ELK Stack May 18 2021 Build mesmerizing visualizations, analytics, and logs from your data using Elasticsearch, Logstash, and Kibana About This Book Solve all your data analytics problems with the ELK stack Explore the power of Kibana4 search and visualizations built over Elasticsearch queries and learn about the features and plugins of Logstash Develop a complete data pipeline using the ELK stack Who This Book Is For If you are a developer or DevOps engineer interested in building a system that provides amazing insights and business metrics out of data sources, of various formats and types, using the open source technology stack that ELK provides, then this book is for you. Basic knowledge of Unix or any programming language will be helpful to make the most out of this book. What You Will Learn Install, configure, and run Elasticsearch, Logstash, and Kibana Understand the need for log analytics and the current challenges in log analysis Build your own data pipeline using the ELK stack Familiarize yourself with the key features of Logstash and the variety of input, filter, and output plugins it provides Build your own custom Logstash plugin Create actionable insights using charts, histograms, and quick search features in Kibana4 Understand the role of Elasticsearch in the ELK stack In Detail The ELK stack—Elasticsearch, Logstash, and Kibana, is a powerful combination of open source tools. Elasticsearch is for deep search and data analytics. Logstash is for centralized logging, log enrichment, and parsing. Kibana is for powerful and beautiful data visualizations. In short, the Elasticsearch ELK stack makes searching and analyzing data easier than ever before. This book will introduce you to the ELK (Elasticsearch, Logstash, and Kibana) stack, starting by showing you how to set up the stack by installing the tools, and basic configuration. You'll move on to building a basic data pipeline using the ELK stack. Next, you'll explore the key features of Logstash and its role in the ELK stack, including creating Logstash plugins, which will enable you to use your own customized plugins. The importance of Elasticsearch and Kibana in the ELK stack is also covered, along with various types of advanced data analysis, and a variety of charts, tables, and maps. Finally, by the end of the book you will be able to develop full-fledged data pipeline using the ELK stack and have a solid understanding of the role of each of the components. Style and approach This book is a step-by-step guide, complete with various examples to solve your data analytics problems by using the ELK stack to explore and visualize data.

Elasticsearch 7 Quick Start Guide Dec 13 2020 Get the most out of Elasticsearch 7's new features to build, deploy, and manage efficient applications Key Features Discover the new features introduced in Elasticsearch 7 Explore techniques for distributed search, indexing, and clustering Gain hands-on knowledge of implementing Elasticsearch for your enterprise Book Description Elasticsearch is one of the most popular tools for distributed search and analytics. This Elasticsearch book highlights the latest features of Elasticsearch 7 and helps you understand how you can use them to build your own search applications with ease. Starting with an introduction to the Elastic Stack, this book will help you quickly get up to speed with using Elasticsearch. You'll learn how to install, configure, manage, secure, and deploy Elasticsearch clusters, as well as how to use your deployment to develop powerful search and analytics solutions. As you progress, you'll also understand how to troubleshoot any issues that you may encounter along the way. Finally, the book will help you explore the inner workings of Elasticsearch and gain insights into queries, analyzers, mappings, and aggregations as you learn to work with search results. By the end of this book, you'll have a basic understanding of how to build and deploy effective search and analytics solutions using Elasticsearch. What you will learn Install Elasticsearch and use it to safely store data and retrieve it when needed Work with a variety of analyzers and filters Discover techniques to improve search results in Elasticsearch Understand how to perform metric and bucket aggregations Implement best practices for moving clusters and applications to production Explore various techniques to secure your Elasticsearch clusters Who this book is for This book is for software developers, engineers, data architects, system administrators, and anyone who wants to get up and running with Elasticsearch 7. No prior experience with Elasticsearch is required.

Sencha Touch in Action Jul 28 2019 Summary Sencha Touch in Action is the definitive guide to developing applications with Sencha Touch. You'll begin with basic design principles for building great mobile applications, and then explore the features of Sencha Touch that bring those ideas to life. You'll learn how and why objects operate in the framework as you work through several real-world examples. This book also promotes the emerging best practices for mobile web development, from widget implementation to developing an application with the Sencha Touch MVC framework. Updated for Sencha Touch 2. About the Technology The Sencha Touch framework makes it easy to build cross-platform mobile apps using HTML5 and JavaScript. It offers numerous features that mimic native mobile APIs and an MVC architecture that feels right at home for application developers. So you get the power and richness of native apps and the convenience of standard web tools. About this Book Sencha Touch in Action is a complete guide for developers of native-quality mobile Sencha Touch applications. You'll explore real-world examples as you master this impressive framework from the ground up. The book shows you good practices for mobile web development, from widget implementation to the structure of MVC applications. The book requires basic JavaScript skills. It assumes no experience with Sencha Touch or Ext JS. What's Inside Covers Sencha Touch 2 Build on your existing web dev skills Create mobile web apps that feel like native apps Extend enterprise apps to mobile clients Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Authors Jay Garcia is a popular speaker, Sencha community advocate, and author of Ext JS in Action. Anthony De Moss is a professional web and mobile developer. Mitchell Simoens is a Sencha developer supporting the Sencha Touch and Ext JS products. Table of Contents PART 1 INTRODUCTION TO SENCHA TOUCH Introducing Sencha Touch Using Sencha Touch for the first time Sencha Touch foundations PART 2 BUILDING MOBILE USER INTERFACES Mastering the building blocks Toolbars, buttons, and docked items Getting the user's attention Data stores and views Working with forms Maps and media PART 3 CONSTRUCTING AN APPLICATION Class system foundations Building Sencha Touch applications

Elasticsearch in Action Nov 04 2022 Elasticsearch makes it easy to add efficient and scalable searches to enterprise applications. Busy administrators and developers love this open source real-time search and analytics engine because they can simply install it, make a few tweaks, and go on with their work. Elasticsearch is miles deep, so once it's up and running, it can be used to build nearly any custom search solution imaginable. Elasticsearch in Action shows how to build scalable search applications using Elasticsearch. It starts off with an informative overview and an engaging introductory example. Within the first few chapters, it discusses core concepts needed to implement basic searches and efficient indexing. With the fundamentals well in hand, readers will gain an organized view of how to optimize their designs. The book focuses on Elasticsearch's REST API via HTTP. Code snippets are written mostly in bash using curl, which makes them easily translatable to other languages. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Mastering Elastic Stack Jul 08 2020 Get the most out of the Elastic Stack for various complex analytics using this comprehensive and practical guide About This Book Your one-stop solution to perform advanced analytics with Elasticsearch, Logstash, and Kibana Learn how to make better sense of your data by searching, analyzing, and logging data in a systematic way This highly practical guide takes you through an advanced implementation on the ELK stack in your enterprise environment Who This Book Is For This book cater to developers using the Elastic stack in their day-to-day work who are familiar with the basics of Elasticsearch, Logstash, and Kibana, and now want to become an expert at using the Elastic stack for data analytics. What You Will Learn Build a pipeline with help of Logstash and Beats to visualize Elasticsearch data in Kibana Use Beats to ship any type of data to the Elastic stack Understand Elasticsearch APIs, modules, and other advanced concepts Explore Logstash

and its plugins Discover how to utilize the new Kibana UI for advanced analytics See how to work with the Elastic Stack using other advanced configurations Customize the Elastic Stack and plugin development for each of the component Work with the Elastic Stack in a production environment Explore the various components of X-Pack in detail. In Detail Even structured data is useless if it can't help you to take strategic decisions and improve existing system. If you love to play with data, or your job requires you to process custom log formats, design a scalable analysis system, and manage logs to do real-time data analysis, this book is your one-stop solution. By combining the massively popular Elasticsearch, Logstash, Beats, and Kibana, elastic.co has advanced the end-to-end stack that delivers actionable insights in real time from almost any type of structured or unstructured data source. If your job requires you to process custom log formats, design a scalable analysis system, explore a variety of data, and manage logs, this book is your one-stop solution. You will learn how to create real-time dashboards and how to manage the life cycle of logs in detail through real-life scenarios. This book brushes up your basic knowledge on implementing the Elastic Stack and then dives deeper into complex and advanced implementations of the Elastic Stack. We'll help you to solve data analytics challenges using the Elastic Stack and provide practical steps on centralized logging and real-time analytics with the Elastic Stack in production. You will get to grip with advanced techniques for log analysis and visualization. Newly announced features such as Beats and X-Pack are also covered in detail with examples. Toward the end, you will see how to use the Elastic stack for real-world case studies and we'll show you some best practices and troubleshooting techniques for the Elastic Stack. Style and approach This practical guide shows you how to perform advanced analytics with the Elastic stack through real-world use cases. It includes common and some not so common scenarios to use the Elastic stack for data analysis.

Architecting and Operating OpenShift Clusters Sep 09 2020 Design and architect resilient OpenShift clusters and gain a keen understanding of how hundreds of projects are integrated into a powerful solution. While there are many OpenShift resources available for developers, this book focuses on the key elements of infrastructure and operations that teams need when looking to integrate and maintain this platform. You'll review important concepts, such as repeatable deployment techniques, advanced OpenShift RBAC capabilities, monitoring clusters, and integrating with external services. You'll also see how to run specialized workloads in OpenShift and how to deploy non-web based applications on the platform, all designed to help cultivate best practices as your organization continue evolve in microservices architectures. OpenShift has become the main enterprise Kubernetes distribution and its market penetration continues to growth at rapid rate. While OpenShift's documentation provides a great list of configuration options to work with the platform, it can be a daunting task to wade through. Architecting and Operating OpenShift Clusters breaks this content down into clear and useful concepts to provide you with a solid understanding of the OpenShift internal architecture. What You'll Learn Operate high availability in multi-tenant OCP clusters Understand OpenShift SDN models, capabilities, and storage classes Integrate OCP with existing data center capabilities and CI/CD pipelines Support advanced capabilities like: Istio, Multus, Kubernetes Operators, hybrid deployments Who This Book Is For Cloud architects, OpenShift cluster administrators, and teams supporting developers in OpenShift environments who have a basic understanding of this platform and microservices architectures.

Relevant Search Aug 01 2022 Summary Relevant Search demystifies relevance work. Using Elasticsearch, it teaches you how to return engaging search results to your users, helping you understand and leverage the internals of Lucene-based search engines. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Users are accustomed to and expect instant, relevant search results. To achieve this, you must master the search engine. Yet for many developers, relevance ranking is mysterious or confusing. About the Book Relevant Search demystifies the subject and shows you that a search engine is a programmable relevance framework. You'll learn how to apply Elasticsearch or Solr to your business's unique ranking problems. The book demonstrates how to program relevance and how to incorporate secondary data sources, taxonomies, text analytics, and personalization. In practice, a relevance framework requires softer skills as well, such as collaborating with stakeholders to discover the right relevance requirements for your business. By the end, you'll be able to achieve a virtuous cycle of provable, measurable relevance improvements over a search product's lifetime. What's Inside Techniques for debugging relevance? Applying search engine features to real problems? Using the user interface to guide searchers? A systematic approach to relevance? A business culture focused on improving search About the Reader For developers trying to build smarter search with Elasticsearch or Solr. About the Authors Doug Turnbull is lead relevance consultant at OpenSource Connections, where he frequently speaks and blogs. John Berryman is a data engineer at Eventbrite, where he specializes in recommendations and search. Foreword author, Trey Grainger, is a director of engineering at CareerBuilder and author of Solr in Action. Table of Contents The search relevance problem Search under the hood Debugging your first relevance problem Taming tokens Basic multifield search Term-centric search Shaping the relevance function Providing relevance feedback Designing a relevance-focused search application The relevance-centered enterprise Semantic and personalized search

Just Spring Oct 30 2019 Get a concise introduction to Spring, the popular open source framework for building lightweight enterprise applications on the Java platform. This example-driven book for Java developers delves into the framework's basic features, as well as complex concepts such as containers. You'll learn how Spring makes Java Messaging Service easier to work with, and how its support for Hibernate helps you work with data persistence and retrieval. In this revised edition of Just Spring, you'll get your hands deep into sample code, beginning with a problem that illustrates Spring's core principle: dependency injection. In the chapters that follow, author Madhusudhan Konda walks you through features that underlie the solution. Dive into the new chapter on advanced concepts, such as bean scopes and property editors Learn dependency injection through a simple object coupling problem Tackle the framework's core fundamentals, including beans and bean factories Discover how Spring makes the Java Messaging Service API easier to use Learn how Spring has revolutionized data access with Java DataBase Connectivity (JDBC) Use Spring with the Hibernate framework to manipulate data as objects

Learning Elasticsearch 7.x Apr 04 2020 A step-by-step guide that will teach you how to use Elasticsearch in your application effectively KEY FEATURES ? Get familiar with the core concepts of Elasticsearch. ? Understand how the search engine works and how Elasticsearch is different from other similar tools. ? Learn to install Elasticsearch on different operating systems. ? Get familiar with the components of Elastic Stack such as Kibana, Logstash, and Beats, etc. ? Learn how to import data from different sources such as RDBMS, and files, etc DESCRIPTION In the modern Information Technology age, we are flooded with loads of data so we should know how to handle those data and transform them to fetch meaningful information. This book is here to help you manage the data using Elasticsearch. The book starts by covering the fundamentals of Elasticsearch and the concept behind it. After the introduction, you will learn how to install Elasticsearch on different platforms. You will then get to know about Index Management where you will learn to create, update, and delete Elasticsearch indices. Then you will understand how the Query DSL works and how to write some complex search queries using the Query DSL. After completing these basic features, you will move to some advanced topics. Under advanced topics, you will learn to handle Geodata which can be used to plot the data on a map. The book then focuses on Data Analysis using Aggregation. You will then learn how to tune Elasticsearch performance. The book ends with a chapter on Elasticsearch administration. WHAT YOU WILL LEARN ? Learn how to create and manage a cluster ? Work with different components of Elastic Stack ? Review the list of top Information Security certifications. ? Get to know more about Elasticsearch Index Management. ? Understand how to improve the performance by tuning Elasticsearch WHO THIS BOOK IS FOR This book is for developers, architects, DBA, DevOps, and other readers who want to learn Elasticsearch efficiently and want to apply that in their application whether it is a new one or an existing one. It is also beneficial to those who want to play with their data using Elasticsearch. Basic computer programming is a prerequisite. TABLE OF CONTENTS 1 Getting started with Elasticsearch 2 Installation Elasticsearch 3 Working with Elastic Stack 4 Preparing your data 5 Importing Data into Elasticsearch 6 Managing Your Index 7 Apply Search on Your Data 8 Handling Geo with Elasticsearch 9 Aggregating Your

Data 10 Improving the Performance 11 Administer Elasticsearch

I Heart Logs Jun 06 2020 Why a book about logs? That's easy: the humble log is an abstraction that lies at the heart of many systems, from NoSQL databases to cryptocurrencies. Even though most engineers don't think much about them, this short book shows you why logs are worthy of your attention. Based on his popular blog posts, LinkedIn principal engineer Jay Kreps shows you how logs work in distributed systems, and then delivers practical applications of these concepts in a variety of common uses—data integration, enterprise architecture, real-time stream processing, data system design, and abstract computing models. Go ahead and take the plunge with logs; you're going to love them. Learn how logs are used for programmatic access in databases and distributed systems Discover solutions to the huge data integration problem when more data of more varieties meet more systems Understand why logs are at the heart of real-time stream processing Learn the role of a log in the internals of online data systems Explore how Jay Kreps applies these ideas to his own work on data infrastructure systems at LinkedIn

Akka in Action Jun 18 2021 Summary Akka in Action is a comprehensive tutorial on building message-oriented systems using Akka. The book takes a hands-on approach, where each new concept is followed by an example that shows you how it works, how to implement the code, and how to (unit) test it. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Akka makes it relatively easy to build applications in the cloud or on devices with many cores that efficiently use the full capacity of the computing power available. It's a toolkit that provides an actor programming model, a runtime, and required support tools for building scalable applications. About the Book Akka in Action shows you how to build message-oriented systems with Akka. This comprehensive, hands-on tutorial introduces each concept with a working example. You'll start with the big picture of how Akka works, and then quickly build and deploy a fully functional REST service out of actors. You'll explore test-driven development and deploying and scaling fault-tolerant systems. After mastering the basics, you'll discover how to model immutable messages, implement domain models, and apply techniques like event sourcing and CQRS. You'll also find a tutorial on building streaming applications using akka-stream and akka-http. Finally, you'll get practical advice on how to customize and extend your Akka system. What's Inside Getting concurrency right Testing and performance tuning Clustered and cloud-based applications Covers Akka version 2.4 About the Reader This book assumes that you're comfortable with Java and Scala. No prior experience with Akka required. About the Authors A software craftsman and architect, Raymond Roestenburg is an Akka committer. Rob Bakker specializes in concurrent back-end systems and systems integration. Rob Williams has more than 20 years of product development experience. Table of Contents Introducing Akka Up and running Test-driven development with actors Fault tolerance Futures Your first distributed Akka app Configuration, logging, and deployment Structural patterns for actors Routing messages Message channels Finite-state machines and agents System integration Streaming Clustering Actor persistence Performance tips Looking ahead

Learning Elasticsearch Sep 21 2021 Store, search, and analyze your data with ease using Elasticsearch 5.x About This Book Get to grips with the basics of Elasticsearch concepts and its APIs, and use them to create efficient applications Create large-scale Elasticsearch clusters and perform analytics using aggregation This comprehensive guide will get you up and running with Elasticsearch 5.x in no time Who This Book Is For If you want to build efficient search and analytics applications using Elasticsearch, this book is for you. It will also benefit developers who have worked with Lucene or Solr before and now want to work with Elasticsearch. No previous knowledge of Elasticsearch is expected. What You Will Learn See how to set up and configure Elasticsearch and Kibana Know how to ingest structured and unstructured data using Elasticsearch Understand how a search engine works and the concepts of relevance and scoring Find out how to query Elasticsearch with a high degree of performance and scalability Improve the user experience by using autocomplete, geolocation queries, and much more See how to slice and dice your data using Elasticsearch aggregations. Grasp how to use Kibana to explore and visualize your data Know how to host on Elastic Cloud and how to use the latest X-Pack features such as Graph and Alerting In Detail Elasticsearch is a modern, fast, distributed, scalable, fault tolerant, and open source search and analytics engine. You can use Elasticsearch for small or large applications with billions of documents. It is built to scale horizontally and can handle both structured and unstructured data. Packed with easy-to-follow examples, this book will ensure you will have a firm understanding of the basics of Elasticsearch and know how to utilize its capabilities efficiently. You will install and set up Elasticsearch and Kibana, and handle documents using the Distributed Document Store. You will see how to query, search, and index your data, and perform aggregation-based analytics with ease. You will see how to use Kibana to explore and visualize your data. Further on, you will learn to handle document relationships, work with geospatial data, and much more, with this easy-to-follow guide. Finally, you will see how you can set up and scale your Elasticsearch clusters in production environments. Style and approach This comprehensive guide will get you started with Elasticsearch 5.x, so you build a solid understanding of the basics. Every topic is explained in depth and is supplemented with practical examples to enhance your understanding.

The Logstash Book Nov 23 2021 Updated for Logstash and ELK 5.0.0! A book designed for SysAdmins, Operations staff, Developers and DevOps who are interested in deploying a log management solution using the open source Elasticsearch, Logstash & Kibana (ELK) stack. In this book we will walk you through installing, deploying, managing and extending Logstash. We're going to do that by introducing you to Example.com, where you're going to start a new job as one of its SysAdmins. The first project you'll be in charge of is developing its new log management solution. We'll teach you how to: * Install and deploy Logstash. * Ship events from a Logstash Shipper to a central Logstash server. * Filter incoming events using a variety of techniques. * Add structured logging to your applications. * Output those events to a selection of useful destinations. * Use Logstash's awesome web interface Kibana. * Scale out your Logstash implementation as your environment grows. * Quickly and easily extend Logstash to deliver additional functionality that you might need. By the end of the book, you should have a functional and effective log management solution that you can deploy into your own environment. Updated for Logstash and ELK 5.0.0!

Spark in Action Dec 25 2021 Summary The Spark distributed data processing platform provides an easy-to-implement tool for ingesting, streaming, and processing data from any source. In Spark in Action, Second Edition, you'll learn to take advantage of Spark's core features and incredible processing speed, with applications including real-time computation, delayed evaluation, and machine learning. Spark skills are a hot commodity in enterprises worldwide, and with Spark's powerful and flexible Java APIs, you can reap all the benefits without first learning Scala or Hadoop. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Analyzing enterprise data starts by reading, filtering, and merging files and streams from many sources. The Spark data processing engine handles this varied volume like a champ, delivering speeds 100 times faster than Hadoop systems. Thanks to SQL support, an intuitive interface, and a straightforward multilanguage API, you can use Spark without learning a complex new ecosystem. About the book Spark in Action, Second Edition, teaches you to create end-to-end analytics applications. In this entirely new book, you'll learn from interesting Java-based examples, including a complete data pipeline for processing NASA satellite data. And you'll discover Java, Python, and Scala code samples hosted on GitHub that you can explore and adapt, plus appendixes that give you a cheat sheet for installing tools and understanding Spark-specific terms. What's inside Writing Spark applications in Java Spark application architecture Ingestion through files, databases, streaming, and Elasticsearch Querying distributed datasets with Spark SQL About the reader This book does not assume previous experience with Spark, Scala, or Hadoop. About the author Jean-Georges Perrin is an experienced data and software architect. He is France's first IBM Champion and has been honored for 12 consecutive years. Table of Contents PART 1 - THE THEORY CRIPPLED BY AWESOME EXAMPLES 1 So, what is Spark, anyway? 2 Architecture and flow 3 The majestic role of the dataframe 4 Fundamentally lazy 5 Building a simple app for deployment 6 Deploying your simple app PART 2 - INGESTION 7 Ingestion from files 8 Ingestion from databases 9 Advanced ingestion: finding data sources and building

your own 10 Ingestion through structured streaming PART 3 - TRANSFORMING YOUR DATA 11 Working with SQL 12 Transforming your data 13 Transforming entire documents 14 Extending transformations with user-defined functions 15 Aggregating your data PART 4 - GOING FURTHER 16 Cache and checkpoint: Enhancing Spark's performances 17 Exporting data and building full data pipelines 18 Exploring deployment

Mastering Kibana 6.x Jan 02 2020 Get to grips with Kibana and its advanced functions to create interactive visualizations and dashboards Key Features Explore visualizations and perform histograms, stats, and map analytics Unleash X-Pack and Timelion, and learn alerting, monitoring, and reporting features Manage dashboards with Beats and create machine learning jobs for faster analytics Book Description Kibana is one of the popular tools among data enthusiasts for slicing and dicing large datasets and uncovering Business Intelligence (BI) with the help of its rich and powerful visualizations. To begin with, Mastering Kibana 6.x quickly introduces you to the features of Kibana 6.x, before teaching you how to create smart dashboards in no time. You will explore metric analytics and graph exploration, followed by understanding how to quickly customize Kibana dashboards. In addition to this, you will learn advanced analytics such as maps, hits, and list analytics. All this will help you enhance your skills in running and comparing multiple queries and filters, influencing your data visualization skills at scale. With Kibana's Timelion feature, you can analyze time series data with histograms and stats analytics. By the end of this book, you will have created a speedy machine learning job using X-Pack capabilities. What you will learn Create unique dashboards with various intuitive data visualizations Visualize Timelion expressions with added histograms and stats analytics Integrate X-Pack with your Elastic Stack in simple steps Extract data from Elasticsearch for advanced analysis and anomaly detection using dashboards Build dashboards from web applications for application logs Create monitoring and alerting dashboards using Beats Who this book is for Mastering Kibana 6.x is for you if you are a big data engineer, DevOps engineer, or data scientist aspiring to go beyond data visualization at scale and gain maximum insights from their large datasets. Basic knowledge of Elasticstack will be an added advantage, although not mandatory.

ElasticSearch Cookbook - Second Edition Apr 16 2021 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.

Microservices in Action May 06 2020 Summary Microservices in Action is a practical book about building and deploying microservice-based applications. Written for developers and architects with a solid grasp of service-oriented development, it tackles the challenge of putting microservices into production. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Invest your time in designing great applications, improving infrastructure, and making the most out of your dev teams. Microservices are easier to write, scale, and maintain than traditional enterprise applications because they're built as a system of independent components. Master a few important new patterns and processes, and you'll be ready to develop, deploy, and run production-quality microservices. About the Book Microservices in Action teaches you how to write and maintain microservice-based applications. Created with day-to-day development in mind, this informative guide immerses you in real-world use cases from design to deployment. You'll discover how microservices enable an efficient continuous delivery pipeline, and explore examples using Kubernetes, Docker, and Google Container Engine. What's inside An overview of microservice architecture Building a delivery pipeline Best practices for designing multi-service transactions and queries Deploying with containers Monitoring your microservices About the Reader Written for intermediate developers familiar with enterprise architecture and cloud platforms like AWS and GCP. About the Author Morgan Bruce and Paulo A. Pereira are experienced engineering leaders. They work daily with microservices in a production environment, using the techniques detailed in this book. Table of Contents PART 1 - The lay of the land Designing and running microservices Microservices at SimpleBank PART 2 - Design Architecture of a microservice application Designing new features Transactions and queries in microservices Designing reliable services Building a reusable microservice framework PART 3 - Deployment Deploying microservices Deployment with containers and schedulers Building a delivery pipeline for microservices PART 4 - Observability and ownership Building a monitoring system Using logs and traces to understand behavior Building microservice teams

Restlet in Action Aug 09 2020 Summary Restlet in Action gets you started with the Restlet Framework and the REST architecture style. You'll create and deploy applications in record time while learning to use popular RESTful Web APIs effectively. This book looks at the many aspects of web development, on both the server and client side, along with cloud computing, mobile Android devices, and Semantic Web applications. About the Technology In a RESTful architecture any component can act, if needed, as both client and server—this is flexible and powerful, but tricky to implement. The Restlet project is a reference implementation with a Java-based API and everything you need to build servers and web clients that integrate with most web and enterprise technologies. About the Book Restlet in Action introduces the Restlet Framework and RESTful web APIs. You'll see how to easily create and deploy your own web API while learning to consume other web APIs effectively. You'll learn about designing, securing, versioning, documentation, optimizing, and more on both the server and client side, as well as about cloud computing, mobile Android devices, and Semantic Web applications. The book requires a basic knowledge of Java and the web, but no prior exposure to REST or Restlet. 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. What's Inside Written by the creators of Restlet! How to create your own web API How to deploy on cloud and mobile platforms Focus on Android, Google App Engine, Google Web Toolkit, and OSGi technologies Table of Contents PART 1 GETTING STARTED Introducing the Restlet Framework Beginning a Restlet application Deploying a Restlet application PART 2 GETTING READY TO ROLL OUT Producing and consuming Restlet representations Securing a Restlet application Documenting and versioning a Restlet application Enhancing a Restlet application with recipes and best practices PART 3 FURTHER USE POSSIBILITIES Using Restlet with cloud platforms Using Restlet in browsers and mobile devices Embracing hypermedia and the Semantic Web The future of Restlet

Lucene in Action Jan 26 2022 When Lucene first hit the scene five years ago, it was nothing short of amazing. By using this open-source, highly scalable, super-fast search engine, developers could integrate search into applications quickly and efficiently. A lot has changed since then—search has grown from a "nice-to-have" feature into an indispensable part of most enterprise applications. Lucene now powers search in diverse companies including Akamai, Netflix, LinkedIn, Technorati, HotJobs, Epiphany, FedEx, Mayo Clinic, MIT, New Scientist Magazine, and many others. Some things remain the same, though. Lucene still delivers high-performance search features in a disarmingly easy-to-use API. Due to its vibrant and diverse open-source community of developers and users, Lucene is relentlessly improving, with evolutions to APIs, significant new features such as payloads, and a huge increase (as much as 8x) in indexing speed with Lucene 2.3. And with clear writing, reusable examples, and unmatched advice on best practices, Lucene in Action, Second Edition is still the definitive guide to developing with 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.

MongoDB in Action Jan 14 2021 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

Elasticsearch in Action, Second Edition Feb 12 2021 Build powerful, production-ready search applications using the incredible features of Elasticsearch. Elasticsearch in Action, Second Edition is a hands-on guide to developing fully functional search engines with Elasticsearch and Kibana. Rewritten for the latest version of Elasticsearch, this completely new second edition explores Elasticsearch's high-level architecture, reveals infrastructure patterns, and walks through the search and analytics capabilities of numerous Elasticsearch APIs. You'll quickly progress from the basics of installation and configuring clusters, to indexing documents, advanced aggregations, and putting your servers into production. By the time you're done, you'll be ready to build amazing search engines for your clients that take advantage of Elasticsearch's modern features. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Solr in Action Apr 28 2022 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

Mastering Search Analytics Aug 28 2019 Readers learn what data to collect, how to analyze it, and how to act upon it. They will learn how to develop everything from an executive level dashboard and ROI measurement to a deep analysis of a specific term or word to see how it can improve overall ranking.

Docker in Action Mar 16 2021 Even small applications have dozens of components. Large applications may have thousands, which makes them challenging to install, maintain, and remove. Docker bundles all application components into a package called a container that keeps things tidy and helps manage any dependencies on other applications or infrastructure. Docker in Action, Second Edition teaches you the skills and knowledge you need to create, deploy, and manage applications hosted in Docker containers. This bestseller has been fully updated with new examples, best practices, and entirely new chapters. You'll start with a clear explanation of the Docker model and learn how to package applications in containers, including techniques for testing and distributing applications. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Storm Applied Oct 11 2020 Summary Storm Applied is a practical guide to using Apache Storm for the real-world tasks associated with processing and analyzing real-time data streams. This immediately useful book starts by building a solid foundation of Storm essentials so that you learn how to think about designing Storm solutions the right way from day one. But it quickly dives into real-world case studies that will bring the novice up to speed with productionizing Storm. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. Summary Storm Applied is a practical guide to using Apache Storm for the real-world tasks associated with processing and analyzing real-time data streams. This immediately useful book starts by building a solid foundation of Storm essentials so that you learn how to think about designing Storm solutions the right way from day one. But it quickly dives into real-world case studies that will bring the novice up to speed with productionizing Storm. About the Technology It's hard to make sense out of data when it's coming at you fast. Like Hadoop, Storm processes large amounts of data but it does it reliably and in real time, guaranteeing that every message will be processed. Storm allows you to scale with your data as it grows, making it an excellent platform to solve your big data problems. About the Book Storm Applied is an example-driven guide to processing and analyzing real-time data streams. This immediately useful book starts by teaching you how to design Storm solutions the right way. Then, it quickly dives into real-world case studies that show you how to scale a high-throughput stream processor, ensure smooth operation within a production cluster, and more. Along the way, you'll learn to use Trident for stateful stream processing, along with other tools from the Storm ecosystem. This book moves through the basics quickly. While prior experience with Storm is not assumed, some experience with big data and real-time systems is helpful. What's Inside Mapping real problems to Storm components Performance tuning and scaling Practical troubleshooting and debugging Exactly-once processing with Trident About the Authors Sean Allen, Matthew Jankowski, and Peter Pathirana lead the development team for a high-volume, search-intensive commercial web application at TheLadders. Table of Contents Introducing Storm Core Storm concepts Topology design Creating robust topologies Moving from local to remote topologies Tuning in Storm Resource contention Storm internals Trident

Getting Started with Kubernetes Dec 01 2019 Learn how to schedule and run application containers using Kubernetes. About This Book Get well-versed with the fundamentals of Kubernetes and get it production-ready for deployments Confidently manage your container clusters and networks using Kubernetes This practical guide will show you container application examples throughout to illustrate the concepts and features of Kubernetes Who This Book Is For This book is for developers, sys admins, and DevOps engineers who want to automate the deployment process and scale their applications. You do not need any knowledge about Kubernetes. What You Will Learn Download, install, and configure the Kubernetes codebase Understand the core concepts of a Kubernetes cluster Be able to set up and access monitoring and logging for Kubernetes clusters Set up external access to applications running in the cluster Understand how CoreOS and Kubernetes can help you achieve greater

performance and container implementation agility Run multiple clusters and manage from a single control plane Explore container security as well as securing Kubernetes clusters Work with third-party extensions and tools In Detail Kubernetes has continued to grow and achieve broad adoption across various industries, helping you to orchestrate and automate container deployments on a massive scale. This book will give you a complete understanding of Kubernetes and how to get a cluster up and running. You will develop an understanding of the installation and configuration process. The book will then focus on the core Kubernetes constructs such as pods, services, replica sets, replication controllers, and labels. You will also understand how cluster level networking is done in Kubernetes. The book will also show you how to manage deployments and perform updates with minimal downtime. Additionally, you will learn about operational aspects of Kubernetes such as monitoring and logging. Advanced concepts such as container security and cluster federation will also be covered. Finally, you will learn about the wider Kubernetes ecosystem with OCP, CoreOS, and Tectonic and explore the third-party extensions and tools that can be used with Kubernetes. By the end of the book, you will have a complete understanding of the Kubernetes platform and will start deploying applications on it. Style and approach This straightforward guide will help you understand how to move your container applications into production through best practices and a step-by-step walkthrough tied to real-world operational strategies.

Learning Kibana 7 Nov 11 2020 A beginner's guide to analyzing and visualizing your Elasticsearch data using Kibana 7 and Timelion Key Features Gain a fundamental understanding of how Kibana operates within the Elastic Stack Explore your data with Elastic Graph and create rich dashboards in Kibana Learn scalable data visualization techniques in Kibana 7 Book Description Kibana is a window into the Elastic Stack, that enables the visual exploration and real-time analysis of your data in Elasticsearch. This book will help you understand the core concepts of the use of Kibana 7 for rich analytics and data visualization. If you're new to the tool or want to get to grips with the latest features introduced in Kibana 7, this book is the perfect beginner's guide. You'll learn how to set up and configure the Elastic Stack and understand where Kibana sits within the architecture. As you advance, you'll learn how to ingest data from different sources using Beats or Logstash into Elasticsearch, followed by exploring and visualizing data in Kibana. Whether working with time series data to create complex graphs using Timelion or embedding visualizations created in Kibana into your web applications, this book covers it all. It also covers topics that every Elastic developer needs to be aware of, such as installing and configuring Application Performance Monitoring (APM) servers and agents. Finally, you'll also learn how to create effective machine learning jobs in Kibana to find anomalies in your data. By the end of this book, you'll have a solid understanding of Kibana, and be able to create your own visual analytics solutions from scratch. What you will learn Explore the data-driven architecture of the Elastic Stack Install and set up Kibana 7 and other Elastic Stack components Use Beats and Logstash to get input from different data sources Create different visualizations using Kibana Build enterprise-grade Elastic dashboards from scratch Use Timelion to play with time series data Install and configure APM servers and APM agents Work with Dev Tools, Spaces, Graph, and other important tools Who this book is for If you're an aspiring Elastic developer or data analysts, this book is for you. You'll also find it useful if you want to get up to speed with the new features of Kibana 7 and perform data visualization on enterprise data. No prior knowledge of Kibana is expected, but some experience with Elasticsearch will be helpful.

Monitoring Elasticsearch Aug 21 2021 Monitor your Elasticsearch cluster's health, and diagnose and solve its performance and reliability issues About This Book Understand common performance and reliability pitfalls in ElasticSearch Use popular monitoring tools such as ElasticSearch-head, BigDesk, Marvel, Kibana, and more This is a step-by-step guide with lots of case studies on solving real-world ElasticSearch cluster issues Who This Book Is For This book is for developers and system administrators who use ElasticSearch in a wide range of capacities. Prior knowledge of ElasticSearch and related technologies would be helpful, but is not necessary. What You Will Learn Explore your cluster with ElasticSearch-head and BigDesk Access the underlying data of the ElasticSearch monitoring plugins using the ElasticSearch API Analyze your cluster's performance with Marvel Troubleshoot some of the common performance and reliability issues that come up when using ElasticSearch Analyze a cluster's historical performance, and get to the bottom of and recover from system failures Use and install various other tools and plugins such as Kibana and Kopf, which is helpful to monitor ElasticSearch In Detail ElasticSearch is a distributed search server similar to Apache Solr with a focus on large datasets, a schema-less setup, and high availability. This schema-free architecture allows ElasticSearch to index and search unstructured content, making it perfectly suited for both small projects and large big data warehouses with petabytes of unstructured data. This book is your toolkit to teach you how to keep your cluster in good health, and show you how to diagnose and treat unexpected issues along the way. You will start by getting introduced to ElasticSearch, and look at some common performance issues that pop up when using the system. You will then see how to install and configure ElasticSearch and the ElasticSearch monitoring plugins. Then, you will proceed to install and use the Marvel dashboard to monitor ElasticSearch. You will find out how to troubleshoot some of the common performance and reliability issues that come up when using ElasticSearch. Finally, you will analyze your cluster's historical performance, and get to know how to get to the bottom of and recover from system failures. This book will guide you through several monitoring tools, and utilizes real-world cases and dilemmas faced when using ElasticSearch, showing you how to solve them simply, quickly, and cleanly. Style and approach This is a step-by-step guide to monitoring your ElasticSearch cluster and correcting performance issues. It is filled with lots of in-depth, real-world use-cases on solving different ElasticSearch cluster issues.

Elasticsearch Server May 30 2022 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.

Neo4j in Action Feb 01 2020 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

Elasticsearch Essentials Oct 23 2021 Harness the power of Elasticsearch to build and manage scalable search and analytics solutions with this fast-paced guide About This Book New to Elasticsearch? Here's what you need—a highly practical guide that gives you a quick start with Elasticsearch using easy-to-follow examples; get up and running with Elasticsearch APIs in no time Get the latest guide on Elasticsearch 2.0.0, which contains concise and adequate information on handling all the issues a developer needs to know while handling data in bulk with search relevancy Learn to create large-scale Elasticsearch clusters using best practices Learn from our experts—written by Bharvi Dixit who has extensive experience in working with search servers (especially Elasticsearch) Who This Book Is For Anyone who wants to build efficient search and analytics applications can choose this book. This book is also beneficial for skilled developers, especially ones experienced with Lucene or Solr, who now want to learn Elasticsearch quickly. What You Will Learn Get to know about advanced Elasticsearch concepts and its REST APIs Write CRUD operations and other search functionalities using the Elasticsearch Python and Java clients Dig into wide range of queries and find out how to use them correctly Design schema and mappings with built-in and custom analyzers Excel in data modeling concepts and query optimization Master document relationships and geospatial data Build analytics using aggregations Setup and scale Elasticsearch clusters using best practices Learn to take data backups and secure Elasticsearch clusters In Detail With constantly evolving and growing datasets, organizations have the need to find actionable insights for their business. Elasticsearch, which is the world's most advanced search and analytics engine, brings the ability to make massive amounts of data usable in a matter of milliseconds. It not only gives you the power to build blazing fast search solutions over a massive amount of data, but can also serve as a NoSQL data store. This guide will take you on a tour to become a competent developer quickly with a solid knowledge level and understanding of the Elasticsearch core concepts. Starting from the beginning, this book will cover these core concepts, setting up Elasticsearch and various plugins, working with analyzers, and creating mappings. This book provides complete coverage of working with Elasticsearch using Python and performing CRUD operations and aggregation-based analytics, handling document relationships in the NoSQL world, working with geospatial data, and taking data backups. Finally, we'll show you how to set up and scale Elasticsearch clusters in production environments as well as providing some best practices. Style and approach This is an easy-to-follow guide with practical examples and clear explanations of the concepts. This fast-paced book believes in providing very rich content focusing majorly on practical implementation. This book will provide you with step-by-step practical examples, letting you know about the common errors and solutions along with ample screenshots and code to ensure your success.

Deep Learning for Search Jun 26 2019 Summary Deep Learning for Search teaches you how to improve the effectiveness of your search by implementing neural network-based techniques. By the time you're finished with the book, you'll be ready to build amazing search engines that deliver the results your users need and that get better as time goes on! Foreword by Chris Mattmann. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning handles the toughest search challenges, including imprecise search terms, badly indexed data, and retrieving images with minimal metadata. And with modern tools like DL4J and TensorFlow, you can apply powerful DL techniques without a deep background in data science or natural language processing (NLP). This book will show you how. About the Book Deep Learning for Search teaches you to improve your search results with neural networks. You'll review how DL relates to search basics like indexing and ranking. Then, you'll walk through in-depth examples to upgrade your search with DL techniques using Apache Lucene and Deeplearning4j. As the book progresses, you'll explore advanced topics like searching through images, translating user queries, and designing search engines that improve as they learn! What's inside Accurate and relevant rankings Searching across languages Content-based image search Search with recommendations About the Reader For developers comfortable with Java or a similar language and search basics. No experience with deep learning or NLP needed. About the Author Tommaso Teofili is a software engineer with a passion for open source and machine learning. As a member of the Apache Software Foundation, he contributes to a number of open source projects, ranging from topics like information retrieval (such as Lucene and Solr) to natural language processing and machine translation (including OpenNLP, Joshua, and UIMA). He currently works at Adobe, developing search and indexing infrastructure components, and researching the areas of natural language processing, information retrieval, and deep learning. He has presented search and machine learning talks at conferences including BerlinBuzzwords, International Conference on Computational Science, ApacheCon, EclipseCon, and others. You can find him on Twitter at @tteofili. Table of Contents PART 1 - SEARCH MEETS DEEP LEARNING Neural search Generating synonyms PART 2 - THROWING NEURAL NETS AT A SEARCH ENGINE From plain retrieval to text generation More-sensitive query suggestions Ranking search results with word embeddings Document embeddings for rankings and recommendations PART 3 - ONE STEP BEYOND Searching across languages Content-based image search A peek at performance

Elasticsearch in Action Oct 03 2022 Summary Elasticsearch in Action teaches you how to build scalable search applications using Elasticsearch. You'll ramp up fast, with an informative overview and an engaging introductory example. Within the first few chapters, you'll pick up the core concepts you need to implement basic searches and efficient indexing. With the fundamentals well in hand, you'll go on to gain an organized view of how to optimize your design. Perfect for developers and administrators building and managing search-oriented applications. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Modern search seems like magic—you type a few words and the search engine appears to know what you want. With the Elasticsearch real-time search and analytics engine, you can give your users this magical experience without having to do complex low-level programming or understand advanced data science algorithms. You just install it, tweak it, and get on with your work. About the Book Elasticsearch in Action teaches you how to write applications that deliver professional quality search. As you read, you'll learn to add basic search features to any application, enhance search results with predictive analysis and relevancy ranking, and use saved data from prior searches to give users a custom experience. This practical book focuses on Elasticsearch's REST API via HTTP. Code snippets are written mostly in bash using cURL, so they're easily translatable to other languages. What's Inside What is a great search application? Building scalable search solutions Using Elasticsearch with any language Configuration and tuning About the Reader For developers and administrators building and managing search-oriented applications. About the Authors Radu Gheorghe is a search consultant and software engineer. Matthew Lee Hinman develops highly available, cloud-based systems. Roy Russo is a specialist in predictive analytics. Table of Contents PART 1 CORE ELASTICSEARCH FUNCTIONALITY Introducing Elasticsearch Diving into the functionality Indexing, updating, and deleting data Searching your data Analyzing your data Searching with relevancy Exploring your data with aggregations Relations among documents PART 2 ADVANCED ELASTICSEARCH FUNCTIONALITY Scaling out Improving performance Administering your cluster

Apache Pulsar in Action Mar 04 2020 Deliver lightning fast and reliable messaging for your distributed applications with the flexible and resilient

Apache Pulsar platform. In Apache Pulsar in Action you will learn how to: Publish from Apache Pulsar into third-party data repositories and platforms Design and develop Apache Pulsar functions Perform interactive SQL queries against data stored in Apache Pulsar Apache Pulsar in Action is a comprehensive and practical guide to building high-traffic applications with Pulsar. You'll learn to use this mature and battle-tested platform to deliver extreme levels of speed and durability to your messaging. Apache Pulsar committer David Kjerrumgaard teaches you to apply Pulsar's seamless scalability through hands-on case studies, including IOT analytics applications and a microservices app based on Pulsar functions. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Reliable server-to-server messaging is the heart of a distributed application. Apache Pulsar is a flexible real-time messaging platform built to run on Kubernetes and deliver the scalability and resilience required for cloud-based systems. Pulsar supports both streaming and message queuing, and unlike other solutions, it can communicate over multiple protocols including MQTT, AMQP, and Kafka's binary protocol. About the book Apache Pulsar in Action teaches you to build scalable streaming messaging systems using Pulsar. You'll start with a rapid introduction to enterprise messaging and discover the unique benefits of Pulsar. Following crystal-clear explanations and engaging examples, you'll use the Pulsar Functions framework to develop a microservices-based application. Real-world case studies illustrate how to implement the most important messaging design patterns. What's inside Publish from Pulsar into third-party data repositories and platforms Design and develop Apache Pulsar functions Create an event-driven food delivery application About the reader Written for experienced Java developers. No prior knowledge of Pulsar required. About the author David Kjerrumgaard is a committer on the Apache Pulsar project. He currently serves as a Developer Advocate for StreamNative, where he develops Pulsar best practices and solutions. Table of Contents PART 1 GETTING STARTED WITH APACHE PULSAR 1 Introduction to Apache Pulsar 2 Pulsar concepts and architecture 3 Interacting with Pulsar PART 2 APACHE PULSAR DEVELOPMENT ESSENTIALS 4 Pulsar functions 5 Pulsar IO connectors 6 Pulsar security 7 Schema registry PART 3 HANDS-ON APPLICATION DEVELOPMENT WITH APACHE PULSAR 8 Pulsar Functions patterns 9 Resiliency patterns 10 Data access 11 Machine learning in Pulsar 12 Edge analytics

Advanced Elasticsearch 7.0 Feb 24 2022 Master the intricacies of Elasticsearch 7.0 and use it to create flexible and scalable search solutions Key Features Master the latest distributed search and analytics capabilities of Elasticsearch 7.0 Perform searching, indexing, and aggregation of your data at scale Discover tips and techniques for speeding up your search query performance Book Description Building enterprise-grade distributed applications and executing systematic search operations call for a strong understanding of Elasticsearch and expertise in using its core APIs and latest features. This book will help you master the advanced functionalities of Elasticsearch and understand how you can develop a sophisticated, real-time search engine confidently. In addition to this, you'll also learn to run machine learning jobs in Elasticsearch to speed up routine tasks. You'll get started by learning to use Elasticsearch features on Hadoop and Spark and make search results faster, thereby improving the speed of query results and enhancing the customer experience. You'll then get up to speed with performing analytics by building a metrics pipeline, defining queries, and using Kibana for intuitive visualizations that help provide decision-makers with better insights. The book will later guide you through using Logstash with examples to collect, parse, and enrich logs before indexing them in Elasticsearch. By the end of this book, you will have comprehensive knowledge of advanced topics such as Apache Spark support, machine learning using Elasticsearch and scikit-learn, and real-time analytics, along with the expertise you need to increase business productivity, perform analytics, and get the very best out of Elasticsearch. What you will learn Pre-process documents before indexing in ingest pipelines Learn how to model your data in the real world Get to grips with using Elasticsearch for exploratory data analysis Understand how to build analytics and RESTful services Use Kibana, Logstash, and Beats for dashboard applications Get up to speed with Spark and Elasticsearch for real-time analytics Explore the basics of Spring Data Elasticsearch, and understand how to index, search, and query in a Spring application Who this book is for This book is for Elasticsearch developers and data engineers who want to take their basic knowledge of Elasticsearch to the next level and use it to build enterprise-grade distributed search applications. Prior experience of working with Elasticsearch will be useful to get the most out of this book.

Elasticsearch 7.0 Cookbook Jul 20 2021 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.