

solution is never one size fits all. Returning the right search results requires conveying domain knowledge and business rules in the search engine's data structures, text analytics, and results ranking capabilities. Relevant Search demystifies relevance work. Using Elasticsearch, it tells how to return engaging search results to users, helping readers understand and leverage the internals of Lucene-based search engines. The book walks through several real-world problems using a cohesive philosophy that combines text analysis, query building, and score shaping to express business ranking rules to the search engine. It outlines how to guide the engineering process by monitoring search user behavior and shifting the enterprise to a search-first culture focused on humans, not computers. It also shows how the search engine provides a deeply pluggable platform for integrating search ranking with machine learning, ontologies, personalization, domain-specific expertise, and other enriching sources. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

C?C++????

?CPython?????,?C????,????????Python????????????Python????????,????????????Python????Python????????.

Over 170 advanced recipes to search, analyze, deploy, manage, and monitor data effectively with Elasticsearch 5.x
About This Book Deploy and manage simple Elasticsearch nodes as well as complex cluster topologies Write native plugins to extend the functionalities of Elasticsearch 5.x to boost your business Packed with clear, step-by-step recipes to walk you through the capabilities of Elasticsearch 5.x Who This Book Is For If you are a developer who wants to get the most out of Elasticsearch for advanced search and analytics, this is the book for you. Some understanding of JSON is expected. If you want to extend Elasticsearch, understanding of Java and related technologies is also required. What You Will Learn Choose the best Elasticsearch cloud topology to deploy and power it up with external plugins Develop tailored mapping to take full control of index steps Build complex queries through managing indices and documents Optimize search results through executing analytics aggregations Monitor the performance of the cluster and nodes Install Kibana to monitor cluster and extend Kibana for plugins Integrate Elasticsearch in Java, Scala, Python and Big Data applications In Detail Elasticsearch is a Lucene-based distributed search server that allows users to index and search unstructured content with petabytes of data. This book is your one-stop guide to master the complete Elasticsearch ecosystem. We'll guide you through comprehensive recipes on what's new in Elasticsearch 5.x, showing you how to create complex queries and analytics, and perform index mapping, aggregation, and scripting. Further on, you will explore the modules of Cluster and Node monitoring and see ways to back up and restore a snapshot of an index. You will understand how to install Kibana to monitor a cluster and also to extend Kibana for plugins. Finally, you will also see how you can integrate your Java, Scala, Python, and Big Data applications such as Apache Spark and Pig with Elasticsearch, and add enhanced functionalities with custom plugins. By the end of this book, you will have an in-depth knowledge of the

backup and recovery. You will get an understanding of how you can scale your Elasticsearch cluster to contextualize it and improve its performance. We'll also show you how you can create your own analysis plugin in Elasticsearch. By the end of the book, you will have all the knowledge necessary to master Elasticsearch and put it to efficient use. Style and approach This comprehensive guide covers intermediate and advanced concepts in Elasticsearch as well as their implementation. An easy-to-follow approach means you'll be able to master even advanced querying, searching, and administration tasks with ease.

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.

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.

??????“??”???????

????????????????(??)??????

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,

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.

Foreword by Werner Vogels, Vice President and Corporate Technology Officer, Amazon The AWS exam has been updated. Your study guide should be, too. The AWS Certified Developer Official Study Guide—Associate Exam is your ultimate preparation resource for the latest exam! Covering all exam objectives, this invaluable resource puts a team of AWS experts at your side with expert guidance, clear explanations, and the wisdom of experience with AWS best practices. You'll master core services and basic architecture, and equip yourself to develop, deploy, and debug cloud-based applications using AWS. The AWS Developer certification is earned by those who demonstrate the technical knowledge and skill associated with best practices for building secure, reliable cloud-based applications using AWS technology. This book is your official exam prep companion, providing everything you need to know to pass with flying colors. Study the AWS Certified Developer Exam objectives Gain expert insight on core AWS services and best practices Test your understanding of key concepts with challenging chapter questions Access online study tools including electronic flashcards, a searchable glossary, practice exams, and more Cloud computing offers businesses the opportunity to replace up-front capital infrastructure expenses with low, variable costs that scale as they grow. This customized responsiveness has negated the need for far-future infrastructure planning, putting thousands of servers at their disposal as needed—and businesses have responded, propelling AWS to the number-one spot among cloud service providers. Now these businesses need qualified AWS developers, and the AWS certification validates the exact skills and knowledge they're looking for. When you're ready to get serious about your cloud credentials, the AWS Certified Developer Official Study Guide—Associate Exam is the resource you need to pass the exam with flying colors. NOTE: As of October 7, 2019, the accompanying code for hands-on exercises in the book is available for downloading from the secure Resources area in the online test bank. You'll find code for Chapters 1, 2, 11, and 12.

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

Master the art of creating scalable, concurrent, and reactive applications using Akka About This Book This book will help you cure anemic models with domain-driven design We cover major Akka programming concepts such as concurrency, scalability, and reactivity You will learn concepts like Event Sourcing and CQRS via Akka Persistence, Akka Streams, Akka Http as well as Akka Clustering Who This Book Is For If you want to use the Lightbend platform to create highly performant reactive applications, then this book is for you. If you are a Scala developer looking for techniques to use all features of the new Akka release and want to incorporate these solutions in your current or new projects, then this book is for you. Expert Java developers who want to build scalable, concurrent, and reactive application will find this book helpful. What You Will Learn Use Akka actors to enable parallel execution Build out domain-driven design based components like entities and aggregates Respond to command requests on that

aggregate root that affect the internal state Leverage Akka Persistence, protobuf and Cassandra to save the persistent state of you entities Build out complex processing graphs with the Graph Builder DSL Understand the dynamic push/pull nature of backpressure handling within Akka Streams Route HTTP requests to an actor and return a response Deploy actor instances across a set of nodes via ConductR for high availability In Detail For a programmer, writing multi-threaded applications is critical as it is important to break large tasks into smaller ones and run them simultaneously. Akka is a distributed computing toolkit that uses the abstraction of the Actor model, enabling developers to build correct, concurrent, and distributed applications using Java and Scala with ease. The book begins with a quick introduction that simplifies concurrent programming with actors. We then proceed to master all aspects of domain-driven design. We'll teach you how to scale out with Akka Remoting/Clustering. Finally, we introduce Conductr as a means to deploy to and manage microservices across a cluster. Style and approach This comprehensive, fast-paced guide is packed with several real-world use cases that will help you understand concepts, issues, and resolutions while using Akka to create highly performant, scalable, and concurrency-proof reactive applications.

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.

???????EJB 3????????????,??EJB 3?????EJB 3?????????Java???API?EJB 3????????????????,???

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

Connect to standard relational databases, noSQL data sources, and more Follow case studies with example components to create your own systems Who This Book Is For Software engineers, architects, and data scientists with an interest in the design and implementation of big data analytical systems using Hadoop, the Hadoop ecosystem, and other associated technologies.

?????Java ?????????????12????????????????????????????????

Exploit the visualization capabilities of Kibana and build powerful interactive dashboards About This Book Introduction to data-driven architecture and the Elastic stack Build effective dashboards for data visualization and explore datasets with Elastic Graph A comprehensive guide to learning scalable data visualization techniques in Kibana Who This Book Is For If you are a developer, data visualization engineer, or data scientist who wants to get the best of data visualization at scale then this book is perfect for you. A basic understanding of Elasticsearch and Logstash is required to make the best use of this book. What You Will Learn How to create visualizations in Kibana Ingest log data, structure an Elasticsearch cluster, and create visualization assets in Kibana Embed Kibana visualization on web pages Scaffold, develop, and deploy new Kibana & Timelion customizations Build a metrics dashboard in Timelion based on time series data Use the Graph plugin visualization feature and leverage a graph query Create, implement, package, and deploy a new custom plugin Use PreAlert to solve anomaly detection challenges In Detail Kibana is an open source data visualization platform that allows you to interact with your data through stunning, powerful graphics. Its simple, browser-based interface enables you to quickly create and share dynamic dashboards that display changes to Elasticsearch queries in real time. In this book, you'll learn how to use the Elastic stack on top of a data architecture to visualize data in real time. All data architectures have different requirements and expectations when it comes to visualizing the data, whether it's logging analytics, metrics, business analytics, graph analytics, or scaling them as per your business requirements. This book will help you master Elastic visualization tools and adapt them to the requirements of your project. You will start by learning how to use the basic visualization features of Kibana 5. Then you will be shown how to implement a pure metric analytics architecture and visualize it using Timelion, a very recent and trendy feature of the Elastic stack. You will learn how to correlate data using the brand-new Graph visualization and build relationships between documents. Finally, you will be familiarized with the setup of a Kibana development environment so that you can build a custom Kibana plugin. By the end of this book you will have all the information needed to take your Elastic stack skills to a new level of data visualization. Style and approach This book takes a comprehensive, step-by-step approach to working with the visualization aspects of the Elastic stack. Every concept is presented in a very easy-to-follow manner that shows you both the logic and method of implementation. Real world cases are referenced to highlight how each of the key concepts can be put to practical use.

[Copyright: 04bde9a4f4b5d27343d5cb1f8d5b9c6b](https://www.manning.com/books/elasticsearch-in-action)