

# Where To Download Apache Solr A Practical Approach To Enterprise Search Read Pdf Free

Apache Solr Mastering Apache Solr Mastering Apache Solr (Colored Version) Solr in Action Solr in Action Relevant Search Practical Hadoop Ecosystem Apache Solr PHP Integration Apache Solr 4 Cookbook Mastering Apache Solr 7.x Expert Angular Solr Cookbook - Third Edition Practical Apache Lucene 8 Scaling Apache Solr Apache Solr Enterprise Search Server - Third Edition Building Search Applications Apache Solr 4 Cookbook Apache Solr Search Patterns Scaling Big Data with Hadoop and Solr Practical Machine Learning: Innovations in Recommendation Taming Text Intelligent Systems in Big Data, Semantic Web and Machine Learning Introduction to Information Retrieval Next-Generation Big Data Administrating Solr Lucene in Action Apache Kafka 1.0 Cookbook Spring Data Devops in Practice Mastering Apache Solr 7. X Scaling Big Data with Hadoop and Solr - Second Edition Ruby in Practice Angular 2 Design Patterns and Best Practices Apache Solr High Performance Hadoop Security Elasticsearch: The Definitive Guide Forms of Judgments and Orders in the High Court of Justice and Court of Appeal Practical Machine Learning: Innovations in Recommendation Building Modular Cloud Apps with OSGi Solr Cookbook - Third Edition

*Elasticsearch: The Definitive Guide* Dec 24 2019 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

**Next-Generation Big Data** Jan 05 2021 Utilize this practical and easy-to-follow guide to modernize traditional enterprise data warehouse and business intelligence environments with next-generation big data technologies. Next-Generation Big Data takes a holistic approach, covering the most important aspects of modern enterprise big data. The book covers not only the main technology stack but also the next-generation tools and applications used for big data warehousing, data warehouse optimization, real-time and batch data ingestion and processing, real-time data visualization, big data governance, data wrangling, big data cloud deployments, and distributed in-memory big data computing. Finally, the book has an extensive and detailed coverage of big data case studies from Navistar, Cerner, British Telecom, Shopzilla, Thomson Reuters, and Mastercard. What You'll Learn Install Apache Kudu, Impala, and Spark to modernize enterprise data warehouse and business intelligence environments, complete with real-world, easy-to-follow examples, and practical advice Integrate HBase, Solr, Oracle, SQL Server, MySQL, Flume, Kafka, HDFS, and Amazon S3 with Apache Kudu, Impala, and Spark Use StreamSets, Talend, Pentaho, and CDAP for real-time and batch data ingestion and processing Utilize Trifacta, Alteryx, and Datameer for data wrangling and interactive data processing Turbocharge Spark with Alluxio, a distributed in-memory storage platform Deploy big data in the cloud using Cloudera Director Perform real-time data visualization and time series analysis using Zoomdata, Apache Kudu, Impala, and Spark Understand enterprise big data topics such as big data governance, metadata management, data lineage, impact analysis, and policy enforcement, and how to use Cloudera Navigator to perform common data governance tasks Implement big data use cases such as big data warehousing, data warehouse optimization, Internet of Things, real-time data ingestion and analytics, complex event processing, and scalable predictive modeling Study real-world big data case studies from innovative companies, including Navistar, Cerner, British Telecom, Shopzilla, Thomson Reuters, and Mastercard Who This Book Is For BI and big data warehouse professionals interested in gaining practical and real-world insight into next-generation big data processing and analytics using Apache Kudu, Impala, and Spark; and those who want to learn more about other advanced enterprise topics **Solr Cookbook - Third Edition** Jan 17 2022 This book is for intermediate Solr Developers who are willing to learn and implement Pro-level practices, techniques, and solutions. This edition will specifically appeal to developers who wish to quickly get to grips with the changes and new features of Apache Solr 5.

**Taming Text** Apr 08 2021 Summary Taming Text, winner of the 2013 Jolt Awards for Productivity, is a hands-on, example-driven guide to working with unstructured text in the context of real-world applications. This book explores how to automatically organize text using approaches such as full-text search, proper name recognition, clustering, tagging, information extraction, and summarization. The book guides you through examples illustrating each of these topics, as well as the foundations upon which they are built. About this Book There is so much text in our lives, we are practically drowning in it. Fortunately, there are innovative tools and techniques for managing unstructured information that can throw the smart developer a much-needed lifeline. You'll find them in this book. Taming Text is a practical, example-driven guide to working with text in real applications. This book introduces you to useful techniques like full-text search, proper name recognition, clustering, tagging, information extraction, and summarization. You'll explore real use cases as you systematically absorb the foundations upon which they are built. Written in a clear and concise style, this book avoids jargon, explaining the subject in terms you can understand without a background in statistics or natural language processing. Examples are in Java, but the concepts can be applied in any language. Written for Java developers, the book requires no prior knowledge of GWT. 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. Winner of 2013 Jolt Awards: The Best Books—one of five notable books every serious programmer should read. What's Inside When to use text-taming techniques Important open-source libraries like Solr and Mahout How to build text-processing applications About the Authors Grant Ingersoll is an engineer, speaker, and trainer, a Lucene committer, and a cofounder of the Mahout machine-learning project. Thomas Morton is the primary developer of OpenNLP and Maximum Entropy. Drew Farris is a technology consultant, software developer, and contributor to Mahout, Lucene, and Solr. "Takes the mystery out of very complex processes."—From the Foreword by Liz Liddy, Dean, iSchool, Syracuse University Table of Contents Getting started taming text Foundations of taming text Searching Fuzzy string matching Identifying people, places, and things Clustering text Classification, categorization, and tagging Building an example question answering system Untamed text: exploring the next frontier

**Apache Solr Enterprise Search Server - Third Edition** Oct 14 2021 This book is for developers who want to learn how to get the most out of Solr in their applications, whether you are new to the field, have used Solr but don't know everything, or simply want a good reference. It would be helpful to have some familiarity with basic programming concepts, but no prior experience is required.

*Practical Machine Learning: Innovations in Recommendation* May 09 2021 Building a simple but powerful recommendation system is much easier than you think. Approachable for all levels of expertise, this report explains innovations that make machine learning practical for business production settings—and demonstrates how even a small-scale development team can design an effective large-scale recommendation system. Apache Mahout committers Ted Dunning and Ellen Friedman walk you through a design that relies on careful simplification. You'll learn how to collect the right data, analyze it with an algorithm from the Mahout library, and then easily deploy the recommender using search technology, such as Apache Solr or Elasticsearch. Powerful and effective, this efficient combination does learning offline and delivers rapid response recommendations in real time. Understand the tradeoffs between simple and complex recommenders Collect user data that tracks user actions—rather than their ratings Predict what a user wants based on behavior by others, using Mahout for co-occurrence analysis Use search technology to offer recommendations in real time, complete with item metadata Watch the recommender in action with a music service example Improve your recommender with dithering, multimodal recommendation, and other techniques

*Intelligent Systems in Big Data, Semantic Web and Machine Learning* Mar 07 2021 This book describes important methodologies, tools and techniques from the fields of artificial intelligence, basically those which are based on relevant conceptual and formal development. The coverage is wide, ranging from machine learning to the use of data on the Semantic Web, with many new topics. The contributions are concerned with machine learning, big data, data processing in medicine, similarity processing in ontologies, semantic image analysis, as well as many applications including the use of machine learning techniques for cloud security, artificial intelligence techniques for detecting COVID-19, the Internet of things, etc. The book is meant to be a very important and useful source of information for researchers and doctoral students in data analysis, Semantic Web, big data, machine learning, computer engineering and related disciplines, as well as for postgraduate students who want to integrate the doctoral cycle.

**Mastering Apache Solr 7. X** Jun 29 2020 Accelerate your enterprise search engine and bring relevancy in your search analytics Key Features A practical guide in building expertise with Indexing, Faceting, Clustering and Pagination Master the management and administration of Enterprise Search Applications and services seamlessly Handle multiple data inputs such as JSON, xml, pdf, doc, xls, ppt, csv and much more. Book Description Apache Solr is the only standalone enterprise search server with a REST-like application interface. providing highly scalable, distributed search and index replication for many of the world's largest internet sites. To begin with, you would be introduced to how you perform full text search, multiple filter search, perform dynamic clustering and so on helping you to brush up the basics of Apache Solr. You will also explore the new features and advanced options released in Apache Solr 7.x which will get you numerous performance aspects and making data investigation simpler, easier and powerful. You will learn to build complex queries, extensive filters and how are they compiled in your system to bring relevance in your search tools. You will learn to carry out Solr scoring, elements affecting the document score and how you can optimize or tune the score for the application at hand. You will learn to extract features of documents, writing complex queries in re-ranking the documents. You will also learn advanced options helping you to know what content is indexed and how the extracted content is indexed. Throughout the book, you would go through complex problems with solutions along with varied approaches to tackle your business needs. By the end of this book, you will gain advanced proficiency to build out-of-box smart search solutions for your enterprise demands. What you will learn Design schema using schema API to access data in the database Advance querying and fine-tuning techniques for better performance Get to grips with indexing using Client API Set up a fault tolerant and highly available server with newer distributed capabilities, SolrCloud Explore Apache Tika to upload data with Solr Cell Understand different data operations that can be done while indexing Master advanced querying through Velocity Search UI, faceting and Query Re-ranking, pagination and spatial search Learn to use JavaScript, Python, SolrJ and Ruby for interacting with Solr Who this book is for The book would rightly appeal to developers, software engineers, data engineers and database architects who are building or seeking to build enterprise-wide effective search engines for business intelligence. Prior experience of Apache Solr or Java programming is must to take the best of this book.

**Forms of Judgments and Orders in the High Court of Justice and Court of Appeal** Nov 22 2019

**Mastering Apache Solr (Colored Version)** Oct 26 2022 Topic: In the open source, full-text search community, a leader emerges – Apache Solr. Apache Solr enables you to index and access documents orders of magnitude faster than classical databases and thereby provides a first-class search experience to your end users. Brief Description: Mastering Apache Solr is a practical, hands-on guide containing crisp, relevant, systematically arranged, and progressive chapters. These chapters contain a wealth of information presented in a direct and easy-to-understand manner. This book covers key technical concepts, highlighting Solr's supremacy over classical databases in full-text search, which will help you accelerate your progress in the Solr world. Detailed Description: Mastering Apache Solr starts with an introduction to Apache Solr, its underlying technologies, the main differences between the classical database engines, and gradually moves to more advanced topics like boosting performance. In this book, we will look under the hood of a large number of topics and discuss answers to pertinent questions like why denormalize data, how to import classical databases' data inside Apache Solr, how to serve Solr through five different web servers, how to optimize them to serve Solr even faster. An important and major topic covered in this book is Solr's querying mechanism, which will prove to be a strong ally in our journey through this book. We then look at boosting performance and deploying Solr using several servlet servers. Finally, we cover how to communicate with Solr using different programming languages, before deploying it in a cloud-based environment. Who this book is for: Mastering Apache Solr has been written for developers, programmers, and data specialists who want to take a leap towards the future of full-text storage and search and offer a world-class experience to their users. The reader is expected to have a working knowledge of traditional databases, Linux-based operating systems, and XML configuration files. Style and Approach: Mastering Apache Solr is written lucidly and has a dynamically simple approach. From the first page to the last, the book remains practical and focuses on the most important topics used in the world of Apache Solr without neglecting important theoretical fundamentals that help you build a strong foundation. Conclusion: Mastering Apache Solr will empower you to provide a world-class search experience to your end users through the discovery of the powerful mechanisms presented in this book.

**Lucene in Action** Nov 03 2020 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.

**Practical Apache Lucene 8** Dec 16 2021 Gain a thorough knowledge of Lucene's capabilities and use it to develop your own search applications. This book explores the Java-based, high-performance text search engine library used to build search capabilities in your applications. Starting with the basics of Lucene and searching, you will learn about the types of queries used in it and also take a look at scoring models. Applying this basic knowledge, you will develop a hello world app using basic Lucene queries and explore functions like scoring and document level boosting. Along the way you will also uncover the concepts of partial searching and matching in Lucene and then learn how to integrate geographical information (geospatial data) in Lucene using spatial queries and n-dimensional indexing. This will prepare you to build a location-aware search engine with a representative data set that allows location constraints to be specified during a search. You'll also develop a text classifier using Lucene and Apache Mahout, a popular machine learning framework. After a detailed review of performance benchmarking and common issues associated with it, you'll learn some of the best practices of tuning the performance of your application. By the end of the book you'll be able to build your first Lucene patch, where you will not only write your patch, but also test it and ensure it adheres to community coding standards. What You'll Learn Master the basics of Apache Lucene Utilize different query types in Apache Lucene Explore scoring and document level boosting Integrate geospatial data into your application Who This Book Is For Developers wanting to learn the finer details of Apache Lucene by developing a series of projects with it.

**Apache Solr Search Patterns** Jul 11 2021 This book is for developers who already know how to use Solr and are looking at procuring advanced strategies for improving their search using Solr. This book is also for people who work with analytics to generate graphs and reports using Solr. Moreover, if you are a search architect who is looking forward to scale your search using Solr, this is a must-have book for you. It would be helpful if you are familiar with the Java programming language.

**Introduction to Information Retrieval** Feb 06 2021 Class-tested and coherent, this textbook teaches classical and web information retrieval, including web search and the related areas of text classification and text clustering from basic concepts. It gives an up-to-date treatment of all aspects of the design and implementation of systems for gathering, indexing, and searching documents; methods for evaluating systems; and an introduction to the use of machine learning methods on text collections. All the important ideas are explained using examples and figures, making it perfect for introductory courses in information retrieval for advanced undergraduates and graduate students in computer science. Based on feedback from extensive classroom experience, the book has been carefully structured in order to make teaching more natural and effective. Slides and additional exercises (with solutions for lecturers) are also available through the book's supporting website to help course instructors prepare their lectures.

**Apache Solr High Performance** Feb 24 2020 This book is an easy-to-follow guide, full of hands-on, real-world examples. Each topic is explained and demonstrated in a specific and user-friendly flow, from search optimization using Solr to Deployment of Zookeeper applications. This book is ideal for Apache Solr developers and want to learn different techniques to optimize Solr performance with utmost efficiency, along with effectively troubleshooting the problems that usually occur while trying to boost performance. Familiarity with search servers and database querying is expected.

**Building Search Applications** Sep 13 2021 Lucene, LingPipe, and Gate are popular open source tools to build powerful search applications. Building Search Applications describes functions from Lucene that include indexing, searching, ranking, and spelling correction to build search engines. With this book you will learn to: Extract tokens from text using custom tokenizers and analyzers from Lucene, LingPipe, and Gate. Construct a search engine index with an optional backend database to manage large document collections. Explore the wide range of Lucene queries to search an index, understand the ranking algorithm for a query, and suggest spelling corrections. Find the names of people, places, and other entities in text using LingPipe and Gate. Categorize documents by topic using classifiers and build groups of self-organized documents using clustering algorithms from LingPipe. Create a Web crawler to scan the Web, Intranet, or desktop using Nutch. Track the sentiment of articles published on the Web with LingPipe.

**Apache Solr 4 Cookbook** Aug 12 2021 "Apache Solr 4 Cookbook" is written in a helpful, practical style with numerous hands-on recipes to help you master Apache Solr to get more precise search results and analysis, higher performance, and reliability. This book is for developers who wish to learn how to master Apache Solr 4. This book will specifically appeal to developers who wish to quickly get to grips with the changes and new features of Apache Solr 4. This book is also handy as a practical guide to solving common problems and issues when using Apache Solr.

**Administrating Solr** Dec 04 2020 A fast-paced, example-based guide to learning how to administrate, monitor, and optimize Apache Solr. "Administrating Solr" is for developers and Solr administrators who have a basic knowledge of Solr and who are looking for ways to keep their Solr server healthy and well maintained. A basic working knowledge of Apache Lucene is recommended, but this is not mandatory.

**Solr Cookbook - Third Edition** Aug 20 2019 This book is for intermediate Solr Developers who are willing to learn and implement Pro-level practices, techniques, and solutions. This edition will specifically appeal to developers who wish to quickly get to grips with the changes and new features of Apache Solr 5.

**Scaling Big Data with Hadoop and Solr** Jun 10 2021 This book is a step-by-step tutorial that will enable you to leverage the flexible search functionality of Apache Solr together with the Big Data power of Apache Hadoop. Scaling Big Data with Hadoop and Solr provides guidance to developers who wish to build high-speed enterprise search platforms using Hadoop and Solr. This book is primarily aimed at Java programmers who wish to extend the Hadoop platform to make it run as an enterprise search without any prior knowledge of Apache Hadoop and Solr.

**DevOps in Practice** Jul 31 2020 DevOps is a cultural and professional movement that's trying to break these walls. Focused on automation, collaboration, tool sharing and knowledge sharing, DevOps has been revealing that developers and system engineers have a lot to learn from one another. In this book, Danilo Sato will show you how to implement DevOps and Continuous Delivery practices so as to raise your system's deployment frequency at the same time as increasing the production application's stability and robustness. You will learn how to automate a web application's build and deploy phases and the infrastructure management, how to monitor the system deployed to production, how to evolve and migrate an architecture to the cloud and still get to know several other tools that you can use on your company

**Solr in Action** Aug 24 2022 A comprehensive guide to using the web application, including such topics as text analysis, faceted search, result grouping, multilingual search, advanced geospatial and data operations, and relevancy tuning.

**Apache Solr PHP Integration** May 21 2022 This book is full of step-by-step example-oriented tutorials which will show readers how to integrate Solr in PHP applications using the available libraries, and boost the inherent search facilities that Solr offers. If you are a developer who knows PHP and is interested in integrating search into your applications, this is the book for you. No advanced knowledge of Solr is required. Very basic knowledge of system commands and the command-line interface on both Linux and Windows is required. You should also be familiar with the concept of Web servers.

**Building Modular Cloud Apps with OSGi** Sep 20 2019 If you're an experienced Java developer in the enterprise, this practical, hands-on book shows you how to use OSGi to design, develop, and deploy modular cloud applications. You'll quickly learn how to use OSGi, through concise code examples and a set of best practices derived from the authors' experiences with real-world projects. Through the course of this book, you'll learn to develop modern web applications with tools and techniques such as RESTful Web Services, NoSQL, provisioning, elasticity, Auto Scaling, hotfixes, and automatic failover. Code samples are available from GitHub. Work with dynamic OSGi services to create modular applications Explore the basics of OSGi bundles and modular application design Learn advanced topics, including semantic versioning, integration testing, and configuring components Understand OSGi pitfalls, anti-patterns, and features you should avoid Create a modular architecture for cloud-based web applications Discover how maintainability, extensibility, scalability, and testability are affected by modular design Get a look at various options for creating web applications with a modular approach Interact with persistent storage services, including relational databases and NoSQL Examine alternatives for deploying modular applications to the cloud

**Practical Hadoop Ecosystem** Jun 22 2022 Learn how to use the Apache Hadoop projects, including MapReduce, HDFS, Apache Hive, Apache HBase, Apache Kafka, Apache Mahout, and Apache Solr. From setting up the environment to running sample applications each chapter in this book is a practical tutorial on using an Apache Hadoop ecosystem project. While several books on Apache Hadoop are available, most are based on the main projects, MapReduce and HDFS, and none discusses the other Apache Hadoop ecosystem projects and how they all work together as a cohesive big data development platform. What You Will Learn: Set up the environment in Linux for Hadoop projects using Cloudera Hadoop Distribution

CDH 5 Run a MapReduce job Store data with Apache Hive, and Apache HBase Index data in HDFS with Apache Solr Develop a Kafka messaging system Stream Logs to HDFS with Apache Flume Transfer data from MySQL database to Hive, HDFS, and HBase with Sqoop Create a Hive table over Apache Solr Develop a Mahout User Recommender System Who This Book Is For: Apache Hadoop developers. Pre-requisite knowledge of Linux and some knowledge of Hadoop is required.

**Ruby in Practice** Apr 27 2020 Ruby in Practice increases your productivity by showing you specific Ruby techniques you can use in your projects. The book offers detailed strategies for using Ruby in a large-scale environment. You'll see concrete examples of integration, messaging, web development, and databases, all presented in a clear Problem/Solution format. This book won't help you push your deadline back, but it will help you get the job done in less time. Above all, Ruby in Practice is a practical book for developers who want an in depth understanding of the Ruby language and its toolset. The book is divided into three major parts. The first part concentrates on issues that developers face both from within their organizations and from their peers. The authors weigh the costs and benefits of using Ruby in enterprise development projects, the over-riding theme being that every developer should always aim to use the right tool for any particular job. The authors then examine the pros and cons of developing with Ruby, and where appropriate, compare Ruby to other languages, both conceptually and in code. The second part discusses techniques for communication and integration of systems. The authors describe how libraries and techniques are used to facilitate messaging, web development and communication automation. They also focus on how the strategies in the first section can be used with these libraries to make development more flexible and efficient. The last part shows how to manage data and integrate with existing data using Ruby. The authors describe techniques for using existing Ruby libraries that either replace popular Java libraries or can interface with existing data in a different way. They go on to show you how strategies covered in the first section can be used to enable Ruby to interact with existing assets or to build new data systems. 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.

**Mastering Apache Solr** Nov 27 2022 Topic: In the open source, full-text search community, a leader emerges – Apache Solr. Apache Solr enables you to index and access documents orders of magnitude faster than classical databases and thereby provides a first-class search experience to your end users. Brief Description: Mastering Apache Solr is a practical, hands-on guide containing crisp, relevant, systematically arranged, and progressive chapters. These chapters contain a wealth of information presented in a direct and easy-to-understand manner. This book covers key technical concepts, highlighting Solr's supremacy over classical databases in full-text search, which will help you accelerate your progress in the Solr world. Detailed Description: Mastering Apache Solr starts with an introduction to Apache Solr, its underlying technologies, the main differences between the classical database engines, and gradually moves to more advanced topics like boosting performance. In this book, we will look under the hood of a large number of topics and discuss answers to pertinent questions like why denormalize data, how to import classical databases' data inside Apache Solr, how to serve Solr through five different web servers, how to optimize them to serve Solr even faster. An important and major topic covered in this book is Solr's querying mechanism, which will prove to be a strong ally in our journey through this book. We then look at boosting performance and deploying Solr using several servlet servers. Finally, we cover how to communicate with Solr using different programming languages, before deploying it in a cloud-based environment. Who this book is for: Mastering Apache Solr has been written for developers, programmers, and data specialists who want to take a leap towards the future of full-text storage and search and offer a world-class experience to their users. The reader is expected to have a working knowledge of traditional databases, Linux-based operating systems, and XML configuration files. Style and Approach: Mastering Apache Solr is written lucidly and has a dynamically simple approach. From the first page to the last, the book remains practical and focuses on the most important topics used in the world of Apache Solr without neglecting important theoretical fundamentals that help you build a strong foundation. Conclusion: Mastering Apache Solr will empower you to provide a world-class search experience to your end users through the discovery of the powerful mechanisms presented in this book.

*Scaling Big Data with Hadoop and Solr - Second Edition* May 29 2020

**Scaling Apache Solr** Nov 15 2021 This book is a step-by-step guide for readers who would like to learn how to build complete enterprise search solutions, with ample real-world examples and case studies. If you are a developer, designer, or architect who would like to build enterprise search solutions for your customers or organization, but have no prior knowledge of Apache Solr/Lucene technologies, this is the book for you.

**Practical Machine Learning: Innovations in Recommendation** Oct 22 2019 Building a simple but powerful recommendation system is much easier than you think. Approachable for all levels of expertise, this report explains innovations that make machine learning practical for business production settings—and demonstrates how even a small-scale development team can design an effective large-scale recommendation system. Apache Mahout committers Ted Dunning and Ellen Friedman walk you through a design that relies on careful simplification. You'll learn how to collect the right data, analyze it with an algorithm from the Mahout library, and then easily deploy the recommender using search technology, such as Apache Solr or Elasticsearch. Powerful and effective, this efficient combination does learning offline and delivers rapid response recommendations in real time. Understand the tradeoffs between simple and complex recommenders Collect user data that tracks user actions—rather than their ratings Predict what a user wants based on behavior by others, using Mahout for co-occurrence analysis Use search technology to offer recommendations in real time, complete with item metadata Watch the recommender in action with a music service example Improve your recommender with dithering, multimodal recommendation, and other techniques

**Solr in Action** Sep 25 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

**Angular 2 Design Patterns and Best Practices** Mar 27 2020 Make the most of Angular by leveraging design patterns and best practices to build stable and high performing apps Key Features Get to grips with the benefits and applicability of using different design patterns in Angular with the help of real-world examples Identify and prevent common problems, programming errors, and anti-patterns Packed with easy-to-follow examples that can be used to create reusable code and extensible designs Book Description This book is an insightful journey through the most valuable design patterns, and it will provide clear guidance on how to use them effectively in Angular. You will explore some of the best ways to work with Angular and how to use it to meet the stability and performance required in today's web development world. You'll get to know some Angular best practices to improve your productivity and the code base of your application. We will take you on a journey through Angular designs for the real world, using a combination of case studies, design patterns to follow, and anti-patterns to avoid. By the end of the book, you will understand the various features of Angular, and will be able to apply well-known, industry-proven design patterns in your work. What you will learn Understand Angular design patterns and anti-patterns Implement the most useful GoF patterns for Angular Explore some of the most famous navigational patterns for Angular Get to know and implement stability patterns Explore and implement operations patterns Explore the official best practices for Angular Monitor and improve the performance of Angular applications Who this book is for If you want to increase your understanding of Angular and apply it to real-life application development, then this book is for you.

**Relevant Search** Jul 23 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

**Expert Angular** Feb 18 2022 Learn everything you need to build highly scalable, robust web applications using Angular release 4 About This Book Apply best practices and design patterns to achieve higher scalability in your Angular applications Understand the latest features of Angular and create your own components Get acquainted with powerful, advanced techniques in Angular to build professional web applications Who This Book Is For This book is for JavaScript developers with some prior exposure to Angular, at least through basic examples. We assume that you've got working knowledge of HTML, CSS, and JavaScript. What You Will Learn Implement asynchronous programming using Angular Beautify your application with the UI components built to the material design specification Secure your web application from unauthorized users Create complex forms, taking full advantage of 2-way data binding Test your Angular applications using the Jasmine and Protractor frameworks for better efficiency Learn how to integrate Angular with Bootstrap to create compelling web applications Use Angular built-in classes to apply animation in your app In Detail Got some experience of Angular under your belt? Want to learn everything about using advanced features for developing websites? This book is everything you need for the deep understanding of Angular that will set you apart from the developer crowd. Angular has introduced a new way to build applications. Creating complex and rich web applications, with a lighter resource footprint, has never been easier or faster. Angular is now at release 4,

with significant changes through previous versions. This book has been written and tested for Angular release 4. Angular is a mature technology, and you'll likely have applications built with earlier versions. This book starts by showing you best practices and approaches to migrating your existing Angular applications so that you can be immediately up-to-date. You will take an in-depth look at components and see how to control the user journey in your applications by implementing routing and navigation. You will learn how to work with asynchronous programming by using Observables. To easily build applications that look great, you will learn all about template syntax and how to beautify applications with Material Design. Mastering forms and data binding will further speed up your application development time. Learning about managing services and animations will help you to progressively enhance your applications. Next you'll use native directives to integrate Bootstrap with Angular. You will see the best ways to test your application with the leading options such as Jasmine and Protractor. At the end of the book, you'll learn how to apply design patterns in Angular, and see the benefits they will bring to your development. Style and approach This book provides comprehensive coverage of all aspects of development with Angular. You will learn about all the most powerful Angular concepts, with examples and best practices. This book is everything you need for the deep understanding of Angular that will set you apart from the developer crowd.

**Apache Solr** Dec 28 2022 Build an enterprise search engine using Apache Solr: index and search documents; ingest data from varied sources; apply various text processing techniques; utilize different search capabilities; and customize Solr to retrieve the desired results. Apache Solr: A Practical Approach to Enterprise Search explains each essential concept-backed by practical and industry examples--to help you attain expert-level knowledge. The book, which assumes a basic knowledge of Java, starts with an introduction to Solr, followed by steps to setting it up, indexing your first set of documents, and searching them. It then introduces you to information retrieval and its implementation in Apache Solr; this will help you understand your search problem, decide the approach to build an effective solution, and use various metrics to evaluate the results. The book next covers the schema design and techniques to build a text analysis chain for cleansing, normalizing and enriching your documents and addressing different types of search queries. It describes various popular matching techniques which are generally applied to improve the precision and recall of searches. You will learn the end-to-end process of data ingestion from varied sources, metadata extraction, pre-processing and transformation of content, various search components, query parsers and other advanced search capabilities. After covering out-of-the-box features, Solr expert Dikshant Shahi dives into ways you can customize Solr for your business and its specific requirements, along with ways to plug in your own components. Most important, you will learn about implementations for Solr scoring, factors affecting the document score, and tuning the score for the application at hand. The book explains why textual scoring is not sufficient for practical ranking of documents and ways to integrate real-world factors for contributing to the document ranking. You'll see how to influence user experience by providing suggestions and recommendations. You'll also see integration of Solr with important related technologies such as OpenNLP and Tika. Additionally, you will learn about scaling Solr using SolrCloud. This book concludes with coverage of semantic search capabilities, which is crucial for taking the search experience to the next level. By the end of Apache Solr, you will be proficient in designing and developing your search engine.

**Apache Solr 4 Cookbook** Apr 20 2022 Over 100 practical recipes to make Apache Solr faster, more reliable and return better results.

**Mastering Apache Solr 7.x** Mar 19 2022 Accelerate your enterprise search engine and bring relevancy in your search analytics Key Features A practical guide in building expertise with Indexing, Faceting, Clustering and Pagination Master the management and administration of Enterprise Search Applications and services seamlessly Handle multiple data inputs such as JSON, xml, pdf, doc, xls,ppt, csv and much more. Book Description Apache Solr is the only standalone enterprise search server with a REST-like application interface. providing highly scalable, distributed search and index replication for many of the world's largest internet sites. To begin with, you would be introduced to how you perform full text search, multiple filter search, perform dynamic clustering and so on helping you to brush up the basics of Apache Solr. You will also explore the new features and advanced options released in Apache Solr 7.x which will get you numerous performance aspects and making data investigation simpler, easier and powerful. You will learn to build complex queries, extensive filters and how are they compiled in your system to bring relevance in your search tools. You will learn to carry out Solr scoring, elements affecting the document score and how you can optimize or tune the score for the application at hand. You will learn to extract features of documents, writing complex queries in re-ranking the documents. You will also learn advanced options helping you to know what content is indexed and how the extracted content is indexed. Throughout the book, you would go through complex problems with solutions along with varied approaches to tackle your business needs. By the end of this book, you will gain advanced proficiency to build out-of-box smart search solutions for your enterprise demands. What you will learn Design schema using schema API to access data in the database Advance querying and fine-tuning techniques for better performance Get to grips with indexing using Client API Set up a fault tolerant and highly available server with newer distributed capabilities, SolrCloud Explore Apache Tika to upload data with Solr Cell Understand different data operations that can be done while indexing Master advanced querying through Velocity Search UI, faceting and Query Re-ranking, pagination and spatial search Learn to use JavaScript, Python, SolrJ and Ruby for interacting with Solr Who this book is for The book would rightly appeal to developers, software engineers, data engineers and database architects who are building or seeking to build enterprise-wide effective search engines for business intelligence. Prior experience of Apache Solr or Java programming is must to take the best of this book.

**Apache Kafka 1.0 Cookbook** Oct 02 2020 Simplify real-time data processing by leveraging the power of Apache Kafka 1.0 Key Features Use Kafka 1.0 features such as Confluent platforms and Kafka streams to build efficient streaming data applications to handle and process your data Integrate Kafka with other Big Data tools such as Apache Hadoop, Apache Spark, and more Hands-on recipes to help you design, operate, maintain, and secure your Apache Kafka cluster with ease Book Description Apache Kafka provides a unified, high-throughput, low-latency platform to handle real-time data feeds. This book will show you how to use Kafka efficiently, and contains practical solutions to the common problems that developers and administrators usually face while working with it. This practical guide contains easy-to-follow recipes to help you set up, configure, and use Apache Kafka in the best possible manner. You will use Apache Kafka Consumers and Producers to build effective real-time streaming applications. The book covers the recently released Kafka version 1.0, the Confluent Platform and Kafka Streams. The programming aspect covered in the book will teach you how to perform important tasks such as message validation, enrichment and composition. Recipes focusing on optimizing the performance of your Kafka cluster, and integrate Kafka with a variety of third-party tools such as Apache Hadoop, Apache Spark, and Elasticsearch will help ease your day to day collaboration with Kafka greatly. Finally, we cover tasks related to monitoring and securing your Apache Kafka cluster using tools such as Ganglia and Graphite. If you're looking to become the go-to person in your organization when it comes to working with Apache Kafka, this book is the only resource you need to have. What you will learn -Install and configure Apache Kafka 1.0 to get optimal performance -Create and configure Kafka Producers and Consumers -Operate your Kafka clusters efficiently by implementing the mirroring technique -Work with the new Confluent platform and Kafka streams, and achieve high availability with Kafka -Monitor Kafka using tools such as Graphite and Ganglia -Integrate Kafka with third-party tools such as Elasticsearch, Logstash, Apache Hadoop, Apache Spark, and more Who this book is for This book is for developers and Kafka administrators who are looking for quick, practical solutions to problems encountered while operating, managing or monitoring Apache Kafka. If you are a developer, some knowledge of Scala or Java will help, while for administrators, some working knowledge of Kafka will be useful.

**Hadoop Security** Jan 25 2020 As more corporations turn to Hadoop to store and process their most valuable data, the risk of a potential breach of those systems increases exponentially. This practical book not only shows Hadoop administrators and security architects how to protect Hadoop data from unauthorized access, it also shows how to limit the ability of an attacker to corrupt or modify data in the event of a security breach. Authors Ben Spivey and Joey Echeverria provide in-depth information about the security features available in Hadoop, and organize them according to common computer security concepts. You'll also get real-world examples that demonstrate how you can apply these concepts to your use cases. Understand the challenges of securing distributed systems, particularly Hadoop Use best practices for preparing Hadoop cluster hardware as securely as possible Get an overview of the Kerberos network authentication protocol Delve into authorization and accounting principles as they apply to Hadoop Learn how to use mechanisms to protect data in a Hadoop cluster, both in transit and at rest Integrate Hadoop data ingest into enterprise-wide security architecture Ensure that security architecture reaches all the way to end-user access

**Spring Data** Sep 01 2020 You can choose several data access frameworks when building Java enterprise applications that work with relational databases. But what about big data? This hands-on introduction shows you how Spring Data makes it relatively easy to build applications across a wide range of new data access technologies such as NoSQL and Hadoop. Through several sample projects, you'll learn how Spring Data provides a consistent programming model that retains NoSQL-specific features and capabilities, and helps you develop Hadoop applications across a wide range of use-cases such as data analysis, event stream processing, and workflow. You'll also discover the features Spring Data adds to Spring's existing JPA and JDBC support for writing RDBMS-based data access layers. Learn about Spring's template helper classes to simplify the use of database-specific functionality Explore Spring Data's repository abstraction and advanced query functionality Use Spring Data with Redis (key/value store), HBase (column-family), MongoDB (document database), and Neo4j (graph database) Discover the GemFire distributed data grid solution Export Spring Data JPA-managed entities to the Web as RESTful web services Simplify the development of HBase applications, using a lightweight object-mapping framework Build example big-data pipelines with Spring Batch and Spring Integration

[artintransit.ca](http://artintransit.ca)