

# Mastering Apache Solr 7x An Expert Guide To Advancing Optimizing And Scaling Your Enterprise Search

Right here, we have countless book **Mastering Apache Solr 7x An Expert Guide To Advancing Optimizing And Scaling Your Enterprise Search** and collections to check out. We additionally provide variant types and also type of the books to browse. The gratifying book, fiction, history, novel, scientific research, as competently as various extra sorts of books are readily comprehensible here.

As this Mastering Apache Solr 7x An Expert Guide To Advancing Optimizing And Scaling Your Enterprise Search, it ends up mammal one of the favored books Mastering Apache Solr 7x An Expert Guide To Advancing Optimizing And Scaling Your Enterprise Search collections that we have. This is why you remain in the best website to see the unbelievable ebook to have.

## **CCNA Security 640-554 Official Cert Guide** - Keith Barker 2013

Trust the best selling Official Cert Guide series from Cisco Press to help you learn, prepare, and practice for exam success. They are built with the objective of providing assessment, review, and practice to help ensure you are fully prepared for your certification exam. CCNA Security 640-554 Official Cert Guide presents you with an organized test preparation routine through the use of proven series elements and techniques. "Do I Know This Already?" quizzes open each chapter and enable you to decide how much time you need to spend on each section. Exam topic lists make referencing easy. Chapter-ending Exam Preparation Tasks help you drill on key concepts you must know thoroughly. · Master Cisco CCNA Security 640-554 exam topics · Assess your knowledge with chapter-opening quizzes · Review key concepts with exam preparation tasks · Practice with realistic exam questions on the CD-ROM CCNA Security 640-554 Official Cert Guide, focuses specifically on the objectives for the Cisco CCNA Security IINS exam. Expert networking professionals Keith Barker and Scott Morris share preparation hints and test-taking tips, helping you identify areas of weakness and improve both your conceptual knowledge and hands-on skills. Material is presented in a concise manner, focusing on increasing your understanding and retention of exam topics. Well-regarded for its level of detail, assessment features, comprehensive design scenarios, and challenging review questions and exercises, this official study guide helps you master the concepts and techniques that will enable you to succeed on the exam the first time. The official study guide helps you master all the topics on the CCNA Security IINS exam, including: Network security concepts Security policies and strategies Network foundation protection (NFP) Cisco Configuration Professional (CCP) Management plane security AAA security Layer 2 security threats IPv6 security Threat mitigation and containment Access Control Lists (ACLs) Network Address Translation (NAT) Cisco IOS zone-based firewalls and ASA firewalls Intrusion prevention and detection systems Public Key Infrastructure (PKI) and cryptography Site-to-site IPsec VPNs and SSL VPNs CCNA Security 640-554 Official Cert Guide is part of a recommended learning path from Cisco that includes simulation and hands-on training from authorized Cisco Learning Partners and self-study products from Cisco Press. To find out more about instructor ...

## **Mastering Apache Solr** - Mr. Mathieu Nayrolles 2014-05-15

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 advance 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.

## **Django 3 By Example** - Antonio Melé 2020-03-31

Learn Django 3 with four end-to-end web projects Key Features Learn Django 3 by building real-world web applications from scratch in Python, using coding best practices Integrate other technologies into your application with clear, step-by-step explanations and comprehensive example code Implement advanced functionalities like a full-text search engine, a user activity stream, or a recommendation engine Add real-time features with Django Channels and WebSockets

**Book Description** If you want to learn the entire process of developing professional web applications with Python and Django, then this book is for you. In the process of building four professional Django projects, you will learn about Django 3 features, how to solve common web development problems, how to implement best practices, and how to successfully deploy your applications. In this book, you will build a blog application, a social image bookmarking website, an online shop, and an e-learning platform. Step-by-step guidance will teach you how to integrate popular technologies, enhance your applications with AJAX, create RESTful APIs, and set up a production environment for your Django projects. By the end of this book, you will have mastered Django 3 by building advanced web applications. What you will learn Build real-world web applications Learn Django essentials, including models, views, ORM, templates, URLs, forms, and authentication Implement advanced features such as custom model fields, custom template tags, cache, middleware, localization, and more Create complex functionalities, such as AJAX interactions, social authentication, a full-text search engine, a payment system, a CMS, a RESTful API, and more Integrate other technologies, including Redis, Celery, RabbitMQ, PostgreSQL, and Channels, into your projects Deploy Django projects in production using NGINX, uWSGI, and Daphne Who this book is for This book is intended for developers with Python knowledge who wish to learn Django in a pragmatic way. Perhaps you are completely new to Django, or you already know a little but you want to get the most out of it. This book will help you to master the most relevant areas of the framework by building practical projects from scratch. You need to have familiarity with programming concepts in order to read this book. Some previous knowledge of HTML and JavaScript is assumed.

**Two Scoops of Django** - Audrey Roy Greenfeld 2015-05-15

"Two scoops of Django introduces you to various tips, tricks, patterns, code snippets, and techniques . . ."--Page 4 of cover.

**Practical Real-time Data Processing and Analytics** - Shilpi Saxena 2017-09-28

A practical guide to help you tackle different real-time data processing and analytics problems using the best tools for each scenario About This Book Learn about the various challenges in real-time data processing and use the right tools to overcome them This book covers popular tools and frameworks such as Spark, Flink, and Apache Storm to solve all your distributed processing problems A practical guide filled with examples, tips, and tricks to help you perform efficient Big Data processing in real-time Who This Book Is For If you are a Java developer who would like to be equipped with all the tools required to devise an end-to-end practical solution on real-time data streaming, then this book is for you. Basic knowledge of real-time processing would be helpful, and knowing the fundamentals of Maven, Shell, and Eclipse would be great. What You Will Learn Get an introduction to the established real-time stack Understand the key integration of all the components Get a thorough understanding of the basic building blocks for real-time solution designing Garnish the search and visualization aspects for your real-time solution Get conceptually and practically acquainted with real-time analytics Be well equipped to apply the knowledge and create your own solutions In Detail With the rise of Big Data, there is an increasing need to process large amounts of data continuously, with a shorter turnaround time. Real-time data processing involves continuous input, processing and output of data, with the condition that the time required for processing is as short as possible. This book covers the majority of the existing and evolving open source technology stack for real-time processing and analytics. You will get to know about all the real-time

solution aspects, from the source to the presentation to persistence. Through this practical book, you'll be equipped with a clear understanding of how to solve challenges on your own. We'll cover topics such as how to set up components, basic executions, integrations, advanced use cases, alerts, and monitoring. You'll be exposed to the popular tools used in real-time processing today such as Apache Spark, Apache Flink, and Storm. Finally, you will put your knowledge to practical use by implementing all of the techniques in the form of a practical, real-world use case. By the end of this book, you will have a solid understanding of all the aspects of real-time data processing and analytics, and will know how to deploy the solutions in production environments in the best possible manner. Style and Approach In this practical guide to real-time analytics, each chapter begins with a basic high-level concept of the topic, followed by a practical, hands-on implementation of each concept, where you can see the working and execution of it. The book is written in a DIY style, with plenty of practical use cases, well-explained code examples, and relevant screenshots and diagrams.

**Data Engineering with Python** - Paul Crickard 2020-10-23

Build, monitor, and manage real-time data pipelines to create data engineering infrastructure efficiently using open-source Apache projects Key Features Become well-versed in data architectures, data preparation, and data optimization skills with the help of practical examples Design data models and learn how to extract, transform, and load (ETL) data using Python Schedule, automate, and monitor complex data pipelines in production Book Description Data engineering provides the foundation for data science and analytics, and forms an important part of all businesses. This book will help you to explore various tools and methods that are used for understanding the data engineering process using Python. The book will show you how to tackle challenges commonly faced in different aspects of data engineering. You'll start with an introduction to the basics of data engineering, along with the technologies and frameworks required to build data pipelines to work with large datasets. You'll learn how to transform and clean data and perform analytics to get the most out of your data. As you advance, you'll discover how to work with big data of varying complexity and production databases, and build data pipelines. Using real-world examples, you'll build architectures on which you'll learn how to deploy data pipelines. By the end of this Python book, you'll have gained a clear understanding of data modeling techniques, and will be able to confidently build data engineering pipelines for tracking data, running quality checks, and making necessary changes in production. What you will learn Understand how data engineering supports data science workflows Discover how to extract data from files and databases and then clean, transform, and enrich it Configure processors for handling different file formats as well as both relational and NoSQL databases Find out how to implement a data pipeline and dashboard to visualize results Use staging and validation to check data before landing in the warehouse Build real-time pipelines with staging areas that perform validation and handle failures Get to grips with deploying pipelines in the production environment Who this book is for This book is for data analysts, ETL developers, and anyone looking to get started with or transition to the field of data engineering or refresh their knowledge of data engineering using Python. This book will also be useful for students planning to build a career in data engineering or IT professionals preparing for a transition. No previous knowledge of data engineering is required.

**Elasticsearch Server** - Rafal Kuc 2013-02-21

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.

*Applying Library Values to Emerging Technology* - Peter D. Fernandez 2018

"Applying Library Values to Emerging Technology: Decision-Making in the Age of Open Access, Maker Spaces, and the Ever-Changing Library offers a wide range of perspectives on how to interpret and apply library values in the context of emerging technologies. Authors include academic librarians, public librarians, and professors, and contributors from the Library Freedom Project, the Institute of Museum and Library Services (IMLS), the Free Ebook Foundation, Creative Commons, the Electronic Frontier Foundation, the Tor Project, the Center for Information Policy Research, and the Institute for the Study of Knowledge Management in Education. Divided into two sections--Contemplating Library Values and Applying Library Values--and using the ALA's Core Values of Librarianship as the primary reference point, chapters emphasize the underlying frameworks that guide librarian practice and capture practical, real-world applications that can ideally serve as a starting point for other librarians encountering similar issues, even if the specific technology or set of values may differ"--website description.

Mastering Apache Solr 7.x - Sandeep Nair 2018-02-22

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.

**Relevant Search** - John Berryman 2016-06-19

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

**Taming Text** - Grant Ingersoll 2012-12-20

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

**Hands-On High Performance with Spring 5** - Chintan Mehta 2018-06-12  
A hands-on guide to creating, monitoring, and tuning a high performance Spring web application Key Features Understand common performance pitfalls and improve your application's performance Build and deploy strategies for complex applications using the microservice architecture Understand internals of JVM - the core of all Java Runtime Environments Book Description While writing an application, performance is paramount. Performance tuning for real-world applications often involves activities geared toward detecting bottlenecks. The recent release of Spring 5.0 brings major advancements in the rich API provided by the Spring framework, which means developers need to master its tools and techniques to achieve high performance applications. Hands-On High Performance with Spring 5 begins with the Spring framework's core features, exploring the integration of different Spring projects. It proceeds to evaluate various Spring specifications to identify those adversely affecting performance. You will learn about bean wiring configurations, aspect-oriented programming, database interaction, and Hibernate to focus on the metrics that help identify performance bottlenecks. You will also look at application monitoring, performance optimization, JVM internals, and garbage collection optimization. Lastly, the book will show you how to leverage the microservice architecture to build a high performance and resilient application. By the end of the book, you will have gained an insight into various techniques and solutions to build and troubleshoot high performance Spring-based applications. What you will learn Master programming best practices and

performance improvement with bean wiring Analyze the performance of various AOP implementations Explore database interactions with Spring to optimize design and configuration Solve Hibernate performance issues and traps Leverage multithreading and concurrent programming to improve application performance Gain a solid foundation in JVM performance tuning using various tools Learn the key concepts of the microservice architecture and how to monitor them Perform Spring Boot performance tuning, monitoring, and health checks Who this book is for If you're a Spring developer who'd like to build high performance applications and have more control over your application's performance in production and development, this book is for you. Some familiarity with Java, Maven, and Eclipse is necessary.

**Apache Solr Enterprise Search Server - Third Edition** - David Smiley 2015-05-26  
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.

**Mastering PostgreSQL 11** - Hans-Jürgen Schönig 2018-10-30  
Master the capabilities of PostgreSQL 11 to efficiently manage and maintain your database Key Features Master advanced concepts of PostgreSQL 11 with real-world datasets and examples Explore query parallelism, data replication, and database performance while working with larger datasets Extend the functionalities of your PostgreSQL instance to suit your organization's needs with minimal effort Book Description This second edition of Mastering PostgreSQL 11 helps you build dynamic database solutions for enterprise applications using the latest release of PostgreSQL, which enables database analysts to design both the physical and technical aspects of the system architecture with ease. This book begins with an introduction to the newly released features in PostgreSQL 11 to help you build efficient and fault-tolerant PostgreSQL applications. You'll examine all of the advanced aspects of PostgreSQL in detail, including logical replication, database clusters, performance tuning, monitoring, and user management. You will also work with the PostgreSQL optimizer, configuring PostgreSQL for high speed, and see how to move from Oracle to PostgreSQL. As you progress through the chapters, you will cover transactions, locking, indexes, and optimizing queries to improve performance. Additionally, you'll learn to manage network security and explore backups and replications, while understanding the useful extensions of PostgreSQL so that you can optimize the speed and performance of large databases. By the end of this book, you will be able to use your database to its utmost capacity by implementing advanced administrative tasks with ease. What you will learn Get to grips with advanced PostgreSQL 11 features and SQL functions Make use of the indexing features in PostgreSQL and fine-tune the performance of your queries Work with stored procedures and manage backup and recovery Master replication and failover techniques Troubleshoot your PostgreSQL instance for solutions to common and not-so-common problems Perform database migration from MySQL and Oracle to PostgreSQL with ease Who this book is for This book is for data and database professionals wanting to implement advanced functionalities and master complex administrative tasks with PostgreSQL 11. Prior experience of database administration with PostgreSQL database will aid in understanding the concepts covered in this book.

**MySQL 8 Administrator's Guide** - Chintan Mehta 2018-02-15  
Step by step guide to monitor, manage, and secure your database engine Key Features Your companion to master all the administration-related tasks in MySQL 8 Ensure high performance and high availability of your MySQL solution using

effective replication and backup techniques A comprehensive guide to performing query optimization, security and a whole host of other administrative tasks in MySQL 8 Book Description MySQL is one of the most popular and widely used relational databases in the world today. The recently released version 8.0 brings along some major advancements in the way your MySQL solution can be administered. This handbook will be your companion to understand the newly introduced features in MySQL and how you can leverage them to design a high-performance MySQL solution for your organization. This book starts with a brief introduction to the newly introduced features in MySQL 8, followed by quickly jumping onto the crucial administration topics that you will find useful in your day to day work. Topics such as migrating to MySQL 8, MySQL benchmarking, achieving high performance by implementing the indexing techniques, and optimizing your queries are covered in this book. You will also learn how to perform replication, scale your MySQL solution and implement effective security techniques. A special section on the common and not so common troubleshooting techniques for effective MySQL administration is also covered in this book. By the end of this highly practical book, you will have all the knowledge you need to tackle any problem you might encounter while administering your MySQL solution. What you will learn Understanding different MySQL 8 data types based on type of contents and storage requirements Best practices for optimal use of features in MySQL 8 Explore globalization configuration and caching techniques to improve performance Create custom storage engine as per system requirements Learn various ways of index implementation for flash memory storages Configure and implement replication along with approaches to use replication as solution Understand how to make your MySQL 8 solution highly available Troubleshoot common issues and identify error codes while using MySQL 8 Who this book is for This book is intended for MySQL administrators who are looking for a handy guide covering all the MySQL administration-related tasks. If you are a DBA looking to get started with MySQL administration, this book will also help you. Knowledge of the basic database concepts is required to get started with this book.

**Biodeterioration of Museum Materials** - S. M. Nair 1983

New York Activity Book - Esther Coombs 2016-09

This fabulously fun New York-based book contains over 20 activities that will entertain children on a day out in the vibrant city of New York. Big pages with plenty of space to draw on and color in will appeal to young children who may still be developing their fine motor skills, while other pages with more detailed drawings and activities will suit older children. Choose from spot the difference in Times Square, doodle pigeons in Central Park, draw an iconic landmark like the Statue of Liberty or Empire State Building or a favorite painting in a gallery. Complete and personalize the pull-out postcards and send them to friends and family. These are just a few of the many activities you'll find that are all centered around the sights of New York. Fun and educational, the New York Activity Book will help children to develop a wide range of skills, including observational, conversational and motor skills, and will make sure they get the most out of their visit to New York.

Solr Cookbook - Third Edition - Rafal Kuc 2015-01-23

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.

Network Security Assessment - Chris McNab 2004

A practical handbook for network administrators who need to develop and implement security assessment programs, exploring a variety of offensive technologies, explaining how to design and deploy networks that are immune to offensive tools and scripts, and detailing an efficient testing model. Original. (Intermediate)

Intelligent Workloads at the Edge - Indraneel Mitra 2022-01-14

Explore IoT, data analytics, and machine learning to solve cyber-physical problems using the latest capabilities of managed services such as AWS IoT Greengrass and Amazon SageMaker Key Features Accelerate your next edge-focused product development with the power of AWS IoT Greengrass Develop proficiency in architecting resilient solutions for the edge with proven best practices Harness the power of analytics and machine learning for solving cyber-physical problems Book Description The Internet of Things (IoT) has transformed how people think about and interact with the world. The ubiquitous deployment of sensors around us makes it possible to study the world at any level of accuracy and enable data-driven decision-making anywhere. Data analytics and machine learning (ML) powered by elastic cloud computing have accelerated our ability to understand and analyze the huge amount of data generated by IoT. Now, edge computing has brought information technologies closer to the data source to lower latency and reduce costs. This book will teach you how to combine the technologies of edge computing, data analytics, and ML to deliver next-generation cyber-physical outcomes. You'll begin by discovering how to create software applications that run on edge devices with AWS IoT Greengrass. As you advance, you'll learn how to process and stream IoT data from the edge to the cloud and use it to train ML models using Amazon SageMaker. The book also shows you how to train these models and run them at the edge for optimized performance, cost savings, and data compliance. By the end of this IoT book, you'll be able to scope your own IoT workloads, bring the power of ML to the edge, and operate those workloads in a production setting. What you will learn Build an end-to-end IoT solution from the edge to the cloud Design and deploy multi-faceted intelligent solutions on the edge Process data at the edge through analytics and ML Package and optimize models for the edge using Amazon SageMaker Implement MLOps and DevOps for operating an edge-based solution Onboard and manage fleets of edge devices at scale Review edge-based workloads against industry best practices Who this book is for This book is for IoT architects and software engineers responsible for delivering analytical and machine learning-backed software solutions to the edge. AWS customers who want to learn and build IoT solutions will find this book useful. Intermediate-level experience with running Python software on Linux is required to make the most of this book.

Mastering Spark with R - Javier Luraschi 2019-10-07

If you're like most R users, you have deep knowledge and love for statistics. But as your organization continues to collect huge amounts of data, adding tools such as Apache Spark makes a lot of sense. With this practical book, data scientists and professionals working with large-scale data applications will learn how to use Spark from R to tackle big data and big compute problems. Authors Javier Luraschi, Kevin Kuo, and Edgar Ruiz show you how to use R with Spark to solve different data analysis problems. This book covers relevant data science topics, cluster computing, and issues that should interest even the most advanced users. Analyze, explore, transform, and visualize data in Apache Spark with R Create statistical models to extract information and predict outcomes; automate the process in production-ready workflows Perform analysis and modeling across many machines using distributed computing techniques Use large-scale data from multiple sources

and different formats with ease from within Spark Learn about alternative modeling frameworks for graph processing, geospatial analysis, and genomics at scale Dive into advanced topics including custom transformations, real-time data processing, and creating custom Spark extensions

*Liferay Beginner's Guide* - Robert Chen 2011-12-01

Part of Packt's Beginner's Guide series, each chapter follows the creation of a fictional neighbourhood site to demonstrate an aspect of Liferay portal with practical examples, screenshots, and step-by-step instructions. All you need in order to benefit from the Liferay Beginner's Guide is programming experience. No prior knowledge of Liferay is required, although experienced Liferay portal programmers who need to get up to speed with its latest features will also find this book useful.

*Ansible for DevOps* - Jeff Geerling 2020-08-05

Ansible is a simple, but powerful, server and configuration management tool. Learn to use Ansible effectively, whether you manage one server--or thousands.

Mastering Elasticsearch - Second Edition - Rafał Kuć 2015-02-27

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.

*Elasticsearch 7 Quick Start Guide* - Anurag Srivastava 2019-10-24

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.

**MySQL 8 for Big Data** - Shabbir Challowala 2017-10-20

Uncover the power of MySQL 8 for Big Data About This Book Combine the powers of MySQL and Hadoop to build a solid Big Data solution for your organization Integrate MySQL with different NoSQL APIs and Big Data tools such as Apache Sqoop A comprehensive guide with practical examples on building a high performance Big Data pipeline with MySQL Who This Book Is For This book is intended for MySQL

database administrators and Big Data professionals looking to integrate MySQL 8 and Hadoop to implement a high performance Big Data solution. Some previous experience with MySQL will be helpful, although the book will highlight the newer features introduced in MySQL 8. What You Will Learn Explore the features of MySQL 8 and how they can be leveraged to handle Big Data Unlock the new features of MySQL 8 for managing structured and unstructured Big Data Integrate MySQL 8 and Hadoop for efficient data processing Perform aggregation using MySQL 8 for optimum data utilization Explore different kinds of join and union in MySQL 8 to process Big Data efficiently Accelerate Big Data processing with Memcached Integrate MySQL with the NoSQL API Implement replication to build highly available solutions for Big Data In Detail With organizations handling large amounts of data on a regular basis, MySQL has become a popular solution to handle this structured Big Data. In this book, you will see how DBAs can use MySQL 8 to handle billions of records, and load and retrieve data with performance comparable or superior to commercial DB solutions with higher costs. Many organizations today depend on MySQL for their websites and a Big Data solution for their data archiving, storage, and analysis needs. However, integrating them can be challenging. This book will show you how to implement a successful Big Data strategy with Apache Hadoop and MySQL 8. It will cover real-time use case scenario to explain integration and achieve Big Data solutions using technologies such as Apache Hadoop, Apache Sqoop, and MySQL Applier. Also, the book includes case studies on Apache Sqoop and real-time event processing. By the end of this book, you will know how to efficiently use MySQL 8 to manage data for your Big Data applications. Style and approach Step by Step guide filled with real-world practical examples.

Web Application Security - Andrew Hoffman 2020-03-02

While many resources for network and IT security are available, detailed knowledge regarding modern web application security has been lacking--until now. This practical guide provides both offensive and defensive security concepts that software engineers can easily learn and apply. Andrew Hoffman, a senior security engineer at Salesforce, introduces three pillars of web application security: recon, offense, and defense. You'll learn methods for effectively researching and analyzing modern web applications--including those you don't have direct access to. You'll also learn how to break into web applications using the latest hacking techniques. Finally, you'll learn how to develop mitigations for use in your own web applications to protect against hackers. Explore common vulnerabilities plaguing today's web applications Learn essential hacking techniques attackers use to exploit applications Map and document web applications for which you don't have direct access Develop and deploy customized exploits that can bypass common defenses Develop and deploy mitigations to protect your applications against hackers Integrate secure coding best practices into your development lifecycle Get practical tips to help you improve the overall security of your web applications

*Apache Solr 4 Cookbook* - Rafał Kuć 2013-01-01

Over 100 practical recipes to make Apache Solr faster, more reliable and return better results.

Solr in Action - Timothy Potter 2014-03-25

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

*MongoDB Applied Design Patterns* - Rick Copeland 2013-03-04

Whether you're building a social media site or an internal-use enterprise application, this hands-on guide shows you the connection between MongoDB and the business problems it's designed to solve. You'll learn how to apply MongoDB design patterns to several challenging domains, such as ecommerce, content management, and online gaming. Using Python and JavaScript code examples, you'll discover how MongoDB lets you scale your data model while simplifying the development process. Many businesses launch NoSQL databases without understanding the techniques for using their features most effectively. This book demonstrates the benefits of document embedding, polymorphic schemas, and other MongoDB patterns for tackling specific big data use cases, including: Operational intelligence: Perform real-time analytics of business data Ecommerce: Use MongoDB as a product catalog master or inventory management system Content management: Learn methods for storing content nodes, binary assets, and discussions Online advertising networks: Apply techniques for frequency capping ad impressions, and keyword targeting and bidding Social networking: Learn how to store a complex social graph, modeled after Google+ Online gaming: Provide concurrent access to character and world data for a multiplayer role-playing game

*The Effective Engineer* - Edmond Lau 2015-03-19

Introducing The Effective Engineer--the only book designed specifically for today's software engineers, based on extensive interviews with engineering leaders at top tech companies, and packed with hundreds of techniques to accelerate your career.

*Learn PostgreSQL* - Luca Ferrari 2020-10-09

A comprehensive guide to building, managing, and securing scalable and reliable database and data warehousing applications using Postgres 12 and 13 Key Features Set up your database cluster and monitor, secure, and fine-tune it for

optimal performance Learn the fundamentals of database management and implement client- and server-side programming using SQL and PL/pgSQL Explore useful tips to develop efficient PostgreSQL database solutions from scratch Book Description PostgreSQL is one of the fastest-growing open source object-relational database management systems (DBMS) in the world. As well as being easy to use, it's scalable and highly efficient. In this book, you'll explore PostgreSQL 12 and 13 and learn how to build database solutions using it. Complete with hands-on tutorials, this guide will teach you how to achieve the right database design required for a reliable environment. You'll learn how to install and configure a PostgreSQL server and even manage users and connections. The book then progresses to key concepts of relational databases, before taking you through the Data Definition Language (DDL) and commonly used DDL commands. To build on your skills, you'll understand how to interact with the live cluster, create database objects, and use tools to connect to the live cluster. You'll then get to grips with creating tables, building indexes, and designing your database schema. Later, you'll explore the Data Manipulation Language (DML) and server-side programming capabilities of PostgreSQL using PL/pgSQL, before learning how to monitor, test, and troubleshoot your database application to ensure high-performance and reliability. By the end of this book, you'll be well-versed with the Postgres database and be able to set up your own PostgreSQL instance and use it to build robust solutions. What you will learn Understand how users and connections are managed by running a PostgreSQL instance Interact with transaction boundaries using server-side programming Identify bottlenecks to maintain your database efficiently Create and manage extensions to add new functionalities to your cluster Choose the best index type for each situation Use online tools to set up a memory configuration that will suit most databases Explore how Postgres can be used in multi-instance environments to provide high-availability, redundancy, and scalability Who this book is for This Postgres book is for anyone interested in learning about the PostgreSQL database from scratch. Anyone looking to build robust data warehousing applications and scale the database for high-availability and performance using the latest features of PostgreSQL will also find this book useful. Although prior knowledge of PostgreSQL is not required, familiarity with databases is expected.

*Using WebPageTest* - Rick Viscomi 2015-10-13

Learn basic and advanced uses of WebPagetest, the performance measurement tool for optimizing websites. This practical guide shows users new to this tool how run tests and interpret results, and helps experienced users gain a better and more thorough understanding of hidden features in WebPagetest that make testing easier. Written by WebPagetest power users and performance experts, this book will help web developers and frontend engineers solve the problem of slow sites. Topics include: Basic test setup--shows beginners how to get meaningful results Advanced test setup--provides another level of technical depth by explaining features not thoroughly documented at webpagetest.org Analysis of results--helps you understand of how to interpret test results Private instance setup--teaches power users the intricacies of the webpagetest private instance and how it works API and external tools--provides a detailed reference for the API and demonstrates tools already using the API to extend WebPagetest

*Amazon SageMaker Best Practices* - Sireesha Muppala 2021-09-24

Overcome advanced challenges in building end-to-end ML solutions by leveraging the capabilities of Amazon SageMaker for developing and integrating ML models into production Key Features Learn best practices for all phases of building machine

learning solutions - from data preparation to monitoring models in productionAutomate end-to-end machine learning workflows with Amazon SageMaker and related AWSDesign, architect, and operate machine learning workloads in the AWS CloudBook Description Amazon SageMaker is a fully managed AWS service that provides the ability to build, train, deploy, and monitor machine learning models. The book begins with a high-level overview of Amazon SageMaker capabilities that map to the various phases of the machine learning process to help set the right foundation. You'll learn efficient tactics to address data science challenges such as processing data at scale, data preparation, connecting to big data pipelines, identifying data bias, running A/B tests, and model explainability using Amazon SageMaker. As you advance, you'll understand how you can tackle the challenge of training at scale, including how to use large data sets while saving costs, monitoring training resources to identify bottlenecks, speeding up long training jobs, and tracking multiple models trained for a common goal. Moving ahead, you'll find out how you can integrate Amazon SageMaker with other AWS to build reliable, cost-optimized, and automated machine learning applications. In addition to this, you'll build ML pipelines integrated with MLOps principles and apply best practices to build secure and performant solutions. By the end of the book, you'll confidently be able to apply Amazon SageMaker's wide range of capabilities to the full spectrum of machine learning workflows. What you will learnPerform data bias detection with AWS Data Wrangler and SageMaker ClarifySpeed up data processing with SageMaker Feature StoreOvercome labeling bias with SageMaker Ground TruthImprove training time with the monitoring and profiling capabilities of SageMaker DebuggerAddress the challenge of model deployment automation with CI/CD using the SageMaker model registryExplore SageMaker Neo for model optimizationImplement data and model quality monitoring with Amazon Model MonitorImprove training time and reduce costs with SageMaker data and model parallelismWho this book is for This book is for expert data scientists responsible for building machine learning applications using Amazon SageMaker.

Working knowledge of Amazon SageMaker, machine learning, deep learning, and experience using Jupyter Notebooks and Python is expected. Basic knowledge of AWS related to data, security, and monitoring will help you make the most of the book.

*PHP & MySQL: Novice to Ninja* - Tom Butler 2017-10-26

PHP & MySQL: Novice to Ninja, 6th Edition is a hands-on guide to learning all the tools, principles, and techniques needed to build a fully functional application using PHP & MySQL. Comprehensively updated to cover PHP 7 and modern best practice, this practical and fun book covers everything from installing PHP and MySQL through to creating a complete online content management system. You'll learn how to: Install PHP & MySQL on Windows, Mac OS X, or Linux Gain a thorough understanding of PHP syntax Use object oriented programming techniques Master database design principles and SQL Develop robust websites that can handle high levels of traffic Build a working content management system (CMS) And much more!

*Liferay Portal Performance Best Practices* - Samir Bhatt 2013-06-10

A step-by-step tutorial on implementing Liferay-based portals to learn performance best practices. The book is good for Liferay portal developers and architects who want to learn performance best practices for implementing Liferay-based solutions. It is assumed that you have a working knowledge of the Liferay portal.

**Python for Finance Cookbook** - Eryk Lewinson 2020-01-31

Solve common and not-so-common financial problems using Python libraries such as NumPy, SciPy, and pandas Key FeaturesUse powerful Python libraries such as pandas,

NumPy, and SciPy to analyze your financial dataExplore unique recipes for financial data analysis and processing with PythonEstimate popular financial models such as CAPM and GARCH using a problem-solution approachBook Description Python is one of the most popular programming languages used in the financial industry, with a huge set of accompanying libraries. In this book, you'll cover different ways of downloading financial data and preparing it for modeling. You'll calculate popular indicators used in technical analysis, such as Bollinger Bands, MACD, RSI, and backtest automatic trading strategies. Next, you'll cover time series analysis and models, such as exponential smoothing, ARIMA, and GARCH (including multivariate specifications), before exploring the popular CAPM and the Fama-French three-factor model. You'll then discover how to optimize asset allocation and use Monte Carlo simulations for tasks such as calculating the price of American options and estimating the Value at Risk (VaR). In later chapters, you'll work through an entire data science project in the financial domain. You'll also learn how to solve the credit card fraud and default problems using advanced classifiers such as random forest, XGBoost, LightGBM, and stacked models. You'll then be able to tune the hyperparameters of the models and handle class imbalance. Finally, you'll focus on learning how to use deep learning (PyTorch) for approaching financial tasks. By the end of this book, you'll have learned how to effectively analyze financial data using a recipe-based approach. What you will learnDownload and preprocess financial data from different sourcesBacktest the performance of automatic trading strategies in a real-world settingEstimate financial econometrics models in Python and interpret their resultsUse Monte Carlo simulations for a variety of tasks such as derivatives valuation and risk assessmentImprove the performance of financial models with the latest Python librariesApply machine learning and deep learning techniques to solve different financial problemsUnderstand the different approaches used to model financial time series dataWho this book is for This book is for financial analysts, data analysts, and Python developers who want to learn how to implement a broad range of tasks in the finance domain. Data scientists looking to devise intelligent financial strategies to perform efficient financial analysis will also find this book useful. Working knowledge of the Python programming language is mandatory to grasp the concepts covered in the book effectively.

**Scaling Big Data with Hadoop and Solr - Second Edition** - Hrishikesh Vijay Karambelkar 2015-04-27

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

*Drupal 9 Module Development* - Daniel Sipos 2020-08-14

Build and extend flexible Drupal sites and applications with this up-to-date, expert guide to Drupal 9 module development Key FeaturesExplore the essential Drupal 9 APIs for module developmentLearn how to implement data modeling, caching, architecture, and much more in your Drupal applicationsDiscover what's new in the latest Drupal core releasesBook Description With its latest release, Drupal 9, the popular open source CMS platform has been updated with new functionalities for building complex Drupal apps with ease. This third edition of the Drupal Module Development guide covers these new Drupal features, helping you to stay on top of code deprecations and the changing architecture with every release. The book starts by introducing you to the Drupal 9 architecture and its subsystems before showing you how to create your first module with basic functionality. You'll



explore the Drupal logging and mailing systems, learn how to output data using the theme layer, and work with menus and links programmatically. Once you've understood the different kinds of data storage, this Drupal guide will demonstrate how to create custom entities and field types and leverage the Database API for lower-level database queries. You'll also learn how to introduce JavaScript into your module, work with various file systems, and ensure that your code works on multilingual sites. Finally, you'll work with Views, create automated tests for your functionality, and write secure code. By the end of the book, you'll have learned how to develop custom modules that can provide solutions to complex business problems, and who knows, maybe you'll even contribute to the Drupal community! What you will learn

Develop custom Drupal 9 modules for your applications  
Master different Drupal 9 subsystems and APIs  
Model, store, manipulate, and process data for effective data management  
Display data and content in a clean and secure way using the theme system  
Test your business logic to prevent regression  
Stay ahead of the curve and write PHP code by implementing best practices

Who this book is for  
If you are a Drupal developer looking to learn Drupal 9 to write modules for your sites, this book is for you. Drupal site builders and PHP developers with basic object-oriented programming skills will also find this book helpful. Although not necessary, some Symfony experience will help with understanding concepts easily.

**Django RESTful Web Services** - Gaston C. Hillar 2018-01-25

Design, build and test RESTful web services with the Django framework and Python  
Key Features  
Create efficient real-world RESTful web services with the latest Django framework  
Authenticate, secure, and integrate third-party packages efficiently in your Web Services  
Leverage the power of Python for faster Web Service development

Book Description  
Django is a Python web framework that makes

the web development process very easy. It reduces the amount of trivial code, which simplifies the creation of web applications and results in faster development. It is very powerful and a great choice for creating RESTful web services. If you are a Python developer and want to efficiently create RESTful web services with Django for your apps, then this is the right book for you. The book starts off by showing you how to install and configure the environment, required software, and tools to create RESTful web services with Django and the Django REST framework. We then move on to working with advanced serialization and migrations to interact with SQLite and non-SQL data sources. We will use the features included in the Django REST framework to improve our simple web service. Further, we will create API views to process diverse HTTP requests on objects, go through relationships and hyperlinked API management, and then discover the necessary steps to include security and permissions related to data models and APIs. We will also apply throttling rules and run tests to check that versioning works as expected. Next we will run automated tests to improve code coverage. By the end of the book, you will be able to build RESTful web services with Django. What you will learn

The best way to build a RESTful Web Service or API with Django and the Django REST Framework  
Develop complex RESTful APIs from scratch with Django and the Django REST Framework  
Work with either SQL or NoSQL data sources  
Design RESTful Web Services based on application requirements  
Use third-party packages and extensions to perform common tasks  
Create automated tests for RESTful web services  
Debug, test, and profile RESTful web services with Django and the Django REST Framework

Who this book is for  
This book is for Python developers who want to create RESTful web services with Django; you need to have a basic working knowledge of Django but no previous experience with RESTful web services is required.