

Redis In Action

If you ally habit such a referred **Redis In Action** ebook that will offer you worth, get the categorically best seller from us currently from several preferred authors. If you desire to humorous books, lots of novels, tale, jokes, and more fictions collections are then launched, from best seller to one of the most current released.

You may not be perplexed to enjoy every ebook collections Redis In Action that we will no question offer. It is not in relation to the costs. Its virtually what you obsession currently. This Redis In Action , as one of the most committed sellers here will unconditionally be among the best options to review.

[Design Patterns](#) - Erich Gamma 1995

Software -- Software Engineering.

Nim in Action - Dominik Picheta 2017-08-04

Summary Nim is a multi-paradigm language that offers powerful customization options with the ability to compile to everything from C to JavaScript. In Nim in Action you'll learn how Nim compares to other languages in style and performance, master its structure and syntax, and discover unique features. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Nim is a multi-paradigm programming language that offers powerful customization options with the ability to compile to everything from C to JavaScript. It can be used in any project and illustrates that you don't have to sacrifice performance for expressiveness!

About the Book Nim in Action is your guide to application development in Nim. You'll learn how Nim compares to other languages in style and performance, master its structure and syntax, and discover unique features. By carefully walking through a Twitter clone and other real-world examples, you'll see just how Nim can be used every day while also learning how to tackle concurrency, package finished applications, and interface with other languages. With the best practices and rich examples in this book, you'll be able to start using Nim today. What's Inside Language features and implementation Nimble package manager Asynchronous I/O Interfacing with C and JavaScript Metaprogramming About the Reader For developers comfortable with mainstream languages like Java, Python, C++ or C#. About the Author Dominik Picheta is one of the principal developers of Nim and author of the

Nimble package manager. Summary PART 1 -THE BASICS OF NIM Why Nim? Getting started PART 2 - NIM IN PRACTICE 3 Writing a chat application 4 A tour through the standard library 5 Package management 6 Parallelism 7 Building a Twitter clone PART 3 - ADVANCED CONCEPTS 8 Interfacing with other languages 9 Metaprogramming

Neo4j in Action - Tareq Abedrabbo 2014-12-05

Summary Neo4j in Action is a comprehensive guide to Neo4j, aimed at application developers and software architects. Using hands-on examples, you'll learn to model graph domains naturally with Neo4j graph structures. The book explores the full power of native Java APIs for graph data manipulation and querying. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Much of the data today is highly connected—from social networks to supply chains to software dependency management—and more connections are continually being uncovered. Neo4j is an ideal graph database tool for highly connected data. It is mature, production-ready, and unique in enabling developers to simply and efficiently model and query connected data. About the Book Neo4j in Action is a comprehensive guide to designing, implementing, and querying graph data using Neo4j. Using hands-on examples, you'll learn to model graph domains naturally with Neo4j graph structures. The book explores the full power of native Java APIs for graph data manipulation and querying. It also covers Cypher, Neo4j's graph query language. Along the way, you'll learn how to integrate Neo4j into your domain-driven app

using Spring Data Neo4j, as well as how to use Neo4j in standalone server or embedded modes. Knowledge of Java basics is required. No prior experience with graph data or Neo4j is assumed. What's Inside Graph database patterns How to model data in social networks How to use Neo4j in your Java applications How to configure and set up Neo4j About the Authors Aleksa Vukotic is an architect specializing in graph data models. Nicki Watt, Dominic Fox, Tareq Abedrabbo, and Jonas Partner work at OpenCredo, a Neo Technology partner, and have been involved in many projects using Neo4j. Table of Contents PART 1 INTRODUCTION TO NEO4J A case for a Neo4j database Data modeling in Neo4j Starting development with Neo4j The power of traversals Indexing the data PART 2 APPLICATION DEVELOPMENT WITH NEO4J Cypher: Neo4j query language Transactions Traversals in depth Spring Data Neo4j PART 3 NEO4J IN PRODUCTION Neo4j: embedded versus server mode [Redis in Action](#) - Josiah Carlson 2013-06-17 Summary Redis in Action introduces Redis and walks you through examples that demonstrate how to use it effectively. You'll begin by getting Redis set up properly and then exploring the key-value model. Then, you'll dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets. Experienced developers will appreciate chapters on clustering and internal scripting to make Redis easier to use. About the Technology When you need near-real-time access to a fast-moving data stream, key-value stores like Redis are the way to go. Redis expands on the key-value pattern by accepting a wide variety of data types, including hashes, strings, lists, and other structures. It provides lightning-fast operations on in-memory datasets, and also makes it easy to persist to disk on the fly. Plus, it's free and open source. About this book Redis in Action introduces Redis and the key-value model. You'll quickly dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets and discover how to integrate with traditional RDBMS or other NoSQL stores. Experienced developers will appreciate the in-depth chapters on clustering and internal scripting. Written for developers familiar with

database concepts. No prior exposure to NoSQL database concepts nor to Redis itself is required. Appropriate for systems administrators comfortable with programming. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside Redis from the ground up Preprocessing real-time data Managing in-memory datasets Pub/sub and configuration Persisting to disk About the Author Dr. Josiah L. Carlson is a seasoned database professional and an active contributor to the Redis community. Table of Contents PART 1 GETTING STARTED Getting to know Redis Anatomy of a Redis web application PART 2 CORE CONCEPTS Commands in Redis Keeping data safe and ensuring performance Using Redis for application support Application components in Redis Search-based applications Building a simple social network PART 3 NEXT STEPS Reducing memory use Scaling Redis Scripting Redis with Lua [Node.js in Action](#) - Tim Oxley 2017-08-16 Summary Node.js in Action, Second Edition is a thoroughly revised book based on the best-selling first edition. It starts at square one and guides you through all the features, techniques, and concepts you'll need to build production-quality Node applications. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology You already know JavaScript. The trick to mastering Node.js is learning how to build applications that fully exploit its powerful asynchronous event handling and non-blocking I/O features. The Node server radically simplifies event-driven real-time apps like chat, games, and live data analytics, and with its incredibly rich ecosystem of modules, tools, and libraries, it's hard to beat! About the Book Based on the bestselling first edition, Node.js in Action, Second Edition is a completely new book. Packed with practical examples, it teaches you how to create high-performance web servers using JavaScript and Node. You'll master key design concepts such as asynchronous programming, state management, and event-driven programming. And you'll learn to put together MVC servers using Express and Connect, design web APIs, and set up the perfect production environment to build, lint, and test. What's Inside Mastering non-blocking I/O The Node event loop Testing and

deploying Web application templating About the Reader Written for web developers with intermediate JavaScript skills. About the Authors The Second Edition author team includes Node masters Alex Young, Bradley Meck, Mike Cantelon, and Tim Oxley, along with original authors Marc Harter, T.J. Holowaychuk, and Nathan Rajlich. Table of contents PART 1 - WELCOME TO NODE Welcome to Node.js Node programming fundamentals What is a Node web application? PART 2 - WEB DEVELOPMENT WITH NODE Front-end build systems Server-side frameworks Connect and Express in depth Web application templating Storing application data Testing Node applications Deploying Node applications and maintaining uptime PART 3 - BEYOND WEB DEVELOPMENT Writing command-line applications Conquering the desktop with Electron

Java Message Service - David A Chappell
2000-12-04

This book is a thorough introduction to Java Message Service (JMS), the standard Java application program interface (API) from Sun Microsystems that supports the formal communication known as "messaging" between computers in a network. JMS provides a common interface to standard messaging protocols and to special messaging services in support of Java programs. The messages exchange crucial data between computers, rather than between users--information such as event notification and service requests. Messaging is often used to coordinate programs in dissimilar systems or written in different programming languages. Using the JMS interface, a programmer can invoke the messaging services of IBM's MQSeries, Progress Software's SonicMQ, and other popular messaging product vendors. In addition, JMS supports messages that contain serialized Java objects and messages that contain Extensible Markup Language (XML) pages. Messaging is a powerful new paradigm that makes it easier to uncouple different parts of an enterprise application. Messaging clients work by sending messages to a message server, which is responsible for delivering the messages to their destination. Message delivery is asynchronous, meaning that the client can continue working without waiting for the message to be delivered. The contents of the

message can be anything from a simple text string to a serialized Java object or an XML document. Java Message Service shows how to build applications using the point-to-point and publish-and-subscribe models; how to use features like transactions and durable subscriptions to make an application reliable; and how to use messaging within Enterprise JavaBeans. It also introduces a new EJB type, the MessageDrivenBean, that is part of EJB 2.0, and discusses integration of messaging into J2EE. *Spring Batch in Action* - Arnaud Cogoluegnes
2011-09-30

Summary Spring Batch in Action is an in-depth guide to writing batch applications using Spring Batch. Written for developers who have basic knowledge of Java and the Spring lightweight container, the book provides both a best-practices approach to writing batch jobs and comprehensive coverage of the Spring Batch framework. About the Technology Even though running batch jobs is a common task, there's no standard way to write them. Spring Batch is a framework for writing batch applications in Java. It includes reusable components and a solid runtime environment, so you don't have to start a new project from scratch. And it uses Spring's familiar programming model to simplify configuration and implementation, so it'll be comfortably familiar to most Java developers. About the Book Spring Batch in Action is a thorough, in-depth guide to writing efficient batch applications. Starting with the basics, it discusses the best practices of batch jobs along with details of the Spring Batch framework. You'll learn by working through dozens of practical, reusable examples in key areas like monitoring, tuning, enterprise integration, and automated testing. No prior batch programming experience is required. Basic knowledge of Java and Spring is assumed. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Batch programming from the ground up Implementing data components Handling errors during batch processing Automating tedious tasks Table of Contents PART 1 BACKGROUND Introducing Spring Batch Spring Batch concepts PART 2 CORE SPRING BATCH Batch configuration Running batch jobs Reading data Writing data Processing data Implementing

bulletproof jobs Transaction management PART 3
ADVANCED SPRING BATCH Controlling execution
Enterprise integration Monitoring jobs Scaling
and parallel processing Testing batch
applications

Spring Microservices in Action - John Carnell
2017-06-11

Summary Spring Microservices in Action teaches you how to build microservice-based applications using Java and the Spring platform. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Microservices break up your code into small, distributed, and independent services that require careful forethought and design. Fortunately, Spring Boot and Spring Cloud simplify your microservice applications, just as the Spring Framework simplifies enterprise Java development. Spring Boot removes the boilerplate code involved with writing a REST-based service. Spring Cloud provides a suite of tools for the discovery, routing, and deployment of microservices to the enterprise and the cloud. About the Book Spring Microservices in Action teaches you how to build microservice-based applications using Java and the Spring platform. You'll learn to do microservice design as you build and deploy your first Spring Cloud application. Throughout the book, carefully selected real-life examples expose microservice-based patterns for configuring, routing, scaling, and deploying your services. You'll see how Spring's intuitive tooling can help augment and refactor existing applications with micro services. What's Inside Core microservice design principles Managing configuration with Spring Cloud Config Client-side resiliency with Spring, Hystrix, and Ribbon Intelligent routing using Netflix Zuul Deploying Spring Cloud applications About the Reader This book is written for developers with Java and Spring experience. About the Author John Carnell is a senior cloud engineer with twenty years of experience in Java. Table of contents Welcome to the cloud, Spring Building microservices with Spring Boot Controlling your configuration with Spring Cloud configuration server On service discovery When bad things happen: client resiliency patterns with Spring Cloud and Netflix Hystrix Service routing with Spring Cloud and Zuul Securing your microservices Event-driven

architecture with Spring Cloud Stream
Distributed tracing with Spring Cloud Sleuth and Zipkin Deploying your microservices
Redis in Action - Josiah L Carlson 2013-06-28
Summary Redis in Action introduces Redis and walks you through examples that demonstrate how to use it effectively. You'll begin by getting Redis set up properly and then exploring the key-value model. Then, you'll dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets. Experienced developers will appreciate chapters on clustering and internal scripting to make Redis easier to use. About the Technology When you need near-real-time access to a fast-moving data stream, key-value stores like Redis are the way to go. Redis expands on the key-value pattern by accepting a wide variety of data types, including hashes, strings, lists, and other structures. It provides lightning-fast operations on in-memory datasets, and also makes it easy to persist to disk on the fly. Plus, it's free and open source. About this book Redis in Action introduces Redis and the key-value model. You'll quickly dive into real use cases including simple caching, distributed ad targeting, and more. You'll learn how to scale Redis from small jobs to massive datasets and discover how to integrate with traditional RDBMS or other NoSQL stores. Experienced developers will appreciate the in-depth chapters on clustering and internal scripting. Written for developers familiar with database concepts. No prior exposure to NoSQL database concepts nor to Redis itself is required. Appropriate for systems administrators comfortable with programming. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside Redis from the ground up Preprocessing real-time data Managing in-memory datasets Pub/sub and configuration Persisting to disk About the Author Dr. Josiah L. Carlson is a seasoned database professional and an active contributor to the Redis community. Table of Contents PART 1 GETTING STARTED Getting to know Redis Anatomy of a Redis web application PART 2 CORE CONCEPTS Commands in Redis Keeping data safe and ensuring performance Using Redis for application support Application components in Redis Search-based

applications Building a simple social network
PART 3 NEXT STEPS Reducing memory use
Scaling Redis Scripting Redis with Lua
[Redis Essentials](#) - Maxwell Dayvson Da Silva
2015-09-08

Harness the power of Redis to integrate and manage your projects efficiently About This Book Learn how to use Redis's data types efficiently to manage large data sets Scale Redis to multiple servers with Twemproxy, Redis Sentinel, and Redis Cluster A fast-paced guide, full of real-world examples to help you get the best out of the features offered by Redis Who This Book Is For If you are a competent developer with experience of working with data structure servers and want to boost your project's performance by learning about features of Redis, then this book is for you. What You Will Learn Build analytics applications using Bitmaps and Hyperloglogs Enhance scalability with Twemproxy, Redis Sentinel, and Redis Cluster Build a Time Series implementation in Node.js and Redis Create your own Redis commands by extending Redis with Lua Get to know security techniques to protect your data (SSL encryption, firewall rules, basic authorization) Persist data to disk and learn the trade-offs of AOF and RDB Understand how to use Node.js, PHP, Python, and Ruby clients for Redis Avoid common pitfalls when designing your next solution In Detail Redis is the most popular in-memory key-value data store. It's very lightweight and its data types give it an edge over the other competitors. If you need an in-memory database or a high-performance cache system that is simple to use and highly scalable, Redis is what you need. Redis Essentials is a fast-paced guide that teaches the fundamentals on data types, explains how to manage data through commands, and shares experiences from big players in the industry. We start off by explaining the basics of Redis followed by the various data types such as Strings, hashes, lists, and more. Next, Common pitfalls for various scenarios are described, followed by solutions to ensure you do not fall into common traps. After this, major differences between client implementations in PHP, Python, and Ruby are presented. Next, you will learn how to extend Redis with Lua, get to know security techniques such as basic authorization, firewall rules, and SSL encryption, and discover how to use

Twemproxy, Redis Sentinel, and Redis Cluster to scale infrastructures horizontally. At the end of this book, you will be able to utilize all the essential features of Redis to optimize your project's performance. Style and approach A practical guide that offers the foundation upon which you can begin to understand the capabilities of Redis using a step-by-step approach. This book is full of real-world problems and in-depth knowledge of the concepts and features of Redis, with plenty of examples.

Professional NoSQL - Shashank Tiwari
2011-08-31

A hands-on guide to leveraging NoSQL databases NoSQL databases are an efficient and powerful tool for storing and manipulating vast quantities of data. Most NoSQL databases scale well as data grows. In addition, they are often malleable and flexible enough to accommodate semi-structured and sparse data sets. This comprehensive hands-on guide presents fundamental concepts and practical solutions for getting you ready to use NoSQL databases. Expert author Shashank Tiwari begins with a helpful introduction on the subject of NoSQL, explains its characteristics and typical uses, and looks at where it fits in the application stack. Unique insights help you choose which NoSQL solutions are best for solving your specific data storage needs. Professional NoSQL: Demystifies the concepts that relate to NoSQL databases, including column-family oriented stores, key/value databases, and document databases. Delves into installing and configuring a number of NoSQL products and the Hadoop family of products. Explains ways of storing, accessing, and querying data in NoSQL databases through examples that use MongoDB, HBase, Cassandra, Redis, CouchDB, Google App Engine Datastore and more. Looks at architecture and internals. Provides guidelines for optimal usage, performance tuning, and scalable configurations. Presents a number of tools and utilities relating to NoSQL, distributed platforms, and scalable processing, including Hive, Pig, RRDtool, Nagios, and more.

[Hadoop in Action](#) - Chuck Lam 2010-11-30

Hadoop in Action teaches readers how to use Hadoop and write MapReduce programs. The intended readers are programmers, architects, and project managers who have to process large amounts of data offline. Hadoop in Action will

lead the reader from obtaining a copy of Hadoop to setting it up in a cluster and writing data analytic programs. The book begins by making the basic idea of Hadoop and MapReduce easier to grasp by applying the default Hadoop installation to a few easy-to-follow tasks, such as analyzing changes in word frequency across a body of documents. The book continues through the basic concepts of MapReduce applications developed using Hadoop, including a close look at framework components, use of Hadoop for a variety of data analysis tasks, and numerous examples of Hadoop in action. Hadoop in Action will explain how to use Hadoop and present design patterns and practices of programming MapReduce. MapReduce is a complex idea both conceptually and in its implementation, and Hadoop users are challenged to learn all the knobs and levers for running Hadoop. This book takes you beyond the mechanics of running Hadoop, teaching you to write meaningful programs in a MapReduce framework. This book assumes the reader will have a basic familiarity with Java, as most code examples will be written in Java. Familiarity with basic statistical concepts (e.g. histogram, correlation) will help the reader appreciate the more advanced data processing examples. 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.

Creating Applications with Mozilla - David Boswell 2002-09-24

Provides guidelines on creating applications with Mozilla that are based on top of the core Mozilla source code. Focuses on utilizing Mozilla's cross-platform development framework.

Web Workers - Ido Green 2012

"Multithreaded programs in JavaScript"--Cover.

Elasticsearch in Action - Roy Russo 2015-11-17

Summary Elasticsearch in Action teaches you how to build scalable search applications using Elasticsearch. You'll ramp up fast, with an informative overview and an engaging introductory example. Within the first few chapters, you'll pick up the core concepts you need to implement basic searches and efficient indexing. With the fundamentals well in hand, you'll go on to gain an organized view of how to optimize your design. Perfect for developers and

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

RabbitMQ in Action - Jason Williams 2012-04-19
Summary RabbitMQ in Action is a fast-paced run through building and managing scalable applications using the RabbitMQ messaging server. It starts by explaining how message queuing works, its history, and how RabbitMQ fits in. Then it shows you real-world examples you can apply to your own scalability and

interoperability challenges. About the Technology There's a virtual switchboard at the core of most large applications where messages race between servers, programs, and services. RabbitMQ is an efficient and easy-to-deploy queue that handles this message traffic effortlessly in all situations, from web startups to massive enterprise systems. About the Book RabbitMQ in Action teaches you to build and manage scalable applications in multiple languages using the RabbitMQ messaging server. It's a snap to get started. You'll learn how message queuing works and how RabbitMQ fits in. Then, you'll explore practical scalability and interoperability issues through many examples. By the end, you'll know how to make Rabbit run like a well-oiled machine in a 24 x 7 x 365 environment. Written for developers familiar with Python, PHP, Java, .NET, or any other modern programming language. No RabbitMQ experience required. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Learn fundamental messaging design patterns Use patterns for on-demand scalability Glue a PHP frontend to a backend written in anything Implement a PubSub-alerting service in 30 minutes flat Configure RabbitMQ's built-in clustering Monitor, manage, extend, and tune RabbitMQ

=====
 ===== Table of Contents
 Pulling RabbitMQ out of the hat Understanding messaging Running and administering Rabbit Solving problems with Rabbit: coding and patterns Clustering and dealing with failure Writing code that survives failure Warrens and Shovels: failover and replication Administering RabbitMQ from the Web Controlling Rabbit with the REST API Monitoring: Houston, we have a problem Supercharging and securing your Rabbit Smart Rabbits: extending RabbitMQ
[HBase in Action](#) - Amandeep Khurana 2012-11-01 Summary HBase in Action has all the knowledge you need to design, build, and run applications using HBase. First, it introduces you to the fundamentals of distributed systems and large scale data handling. Then, you'll explore real-world applications and code samples with just enough theory to understand the practical techniques. You'll see how to build applications

with HBase and take advantage of the MapReduce processing framework. And along the way you'll learn patterns and best practices. About the Technology HBase is a NoSQL storage system designed for fast, random access to large volumes of data. It runs on commodity hardware and scales smoothly from modest datasets to billions of rows and millions of columns. About this Book HBase in Action is an experience-driven guide that shows you how to design, build, and run applications using HBase. First, it introduces you to the fundamentals of handling big data. Then, you'll explore HBase with the help of real applications and code samples and with just enough theory to back up the practical techniques. You'll take advantage of the MapReduce processing framework and benefit from seeing HBase best practices in action. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside When and how to use HBase Practical examples Design patterns for scalable data systems Deployment, integration, and design Written for developers and architects familiar with data storage and processing. No prior knowledge of HBase, Hadoop, or MapReduce is required. Table of Contents PART 1 HBASE FUNDAMENTALS Introducing HBase Getting started Distributed HBase, HDFS, and MapReduce PART 2 ADVANCED CONCEPTS HBase table design Extending HBase with coprocessors Alternative HBase clients PART 3 EXAMPLE APPLICATIONS HBase by example: OpenTSDB Scaling GIS on HBase PART 4 OPERATIONALIZING HBASE Deploying HBase Operations
Ruby on Rails Tutorial - Michael Hartl 2016-11-17 This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. Used by sites as varied as Twitter, GitHub, Disney, and Airbnb, Ruby on Rails is one of the most popular frameworks for developing web applications, but it can be challenging to learn and use. Whether you're new to web development or new only to Rails, Ruby on Rails™ Tutorial, Fourth Edition, is the solution. Best-selling author and leading Rails developer Michael Hartl teaches Rails by guiding you through the development of three example

applications of increasing sophistication. The tutorial's examples focus on the general principles of web development needed for virtually any kind of website. The updates to this edition include full compatibility with Rails 5, a division of the largest chapters into more manageable units, and a huge number of new exercises interspersed in each chapter for maximum reinforcement of the material. This indispensable guide provides integrated tutorials not only for Rails, but also for the essential Ruby, HTML, CSS, and SQL skills you need when developing web applications. Hartl explains how each new technique solves a real-world problem, and then he demonstrates it with bite-sized code that's simple enough to understand, yet novel enough to be useful. Whatever your previous web development experience, this book will guide you to true Rails mastery. This book will help you Install and set up your Rails development environment, including pre-installed integrated development environment (IDE) in the cloud Go beyond generated code to truly understand how to build Rails applications from scratch Learn testing and test-driven development (TDD) Effectively use the Model-View-Controller (MVC) pattern Structure applications using the REST architecture Build static pages and transform them into dynamic ones Master the Ruby programming skills all Rails developers need Create high-quality site layouts and data models Implement registration and authentication systems, including validation and secure passwords Update, display, and delete users Upload images in production using a cloud storage service Implement account activation and password reset, including sending email with Rails Add social features and microblogging, including an introduction to Ajax Record version changes with Git and create a secure remote repository at Bitbucket Deploy your applications early and often with Heroku

Making Sense of NoSQL - Ann Kelly
2013-09-02

Summary Making Sense of NoSQL clearly and concisely explains the concepts, features, benefits, potential, and limitations of NoSQL technologies. Using examples and use cases, illustrations, and plain, jargon-free writing, this guide shows how you can effectively assemble a NoSQL solution to replace or augment the

traditional RDBMS you have now. About this Book If you want to understand and perhaps start using the new data storage and analysis technologies that go beyond the SQL database model, this book is for you. Written in plain language suitable for technical managers and developers, and using many examples, use cases, and illustrations, this book explains the concepts, features, benefits, potential, and limitations of NoSQL. Making Sense of NoSQL starts by comparing familiar database concepts to the new NoSQL patterns that augment or replace them. Then, you'll explore case studies on big data, search, reliability, and business agility that apply these new patterns to today's business problems. You'll see how NoSQL systems can leverage the resources of modern cloud computing and multiple-CPU data centers. The final chapters show you how to choose the right NoSQL technologies for your own needs. Managers and developers will welcome this lucid overview of the potential and capabilities of NoSQL technologies. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside NoSQL data architecture patterns NoSQL for big data Search, high availability, and security Choosing an architecture About the Authors Dan McCreary and Ann Kelly lead an independent training and consultancy firm focused on NoSQL solutions and are cofounders of the NoSQL Now! Conference. Table of Contents PART 1 INTRODUCTION NoSQL: It's about making intelligent choices NoSQL concepts PART 2 DATABASE PATTERNS Foundational data architecture patterns NoSQL data architecture patterns Native XML databases PART 3 NOSQL SOLUTIONS Using NoSQL to manage big data Finding information with NoSQL search Building high-availability solutions with NoSQL Increasing agility with NoSQL PART 4 ADVANCED TOPICS NoSQL and functional programming Security: protecting data in your NoSQL systems Selecting the right NoSQL solution

Redis 4.x Cookbook - Pengcheng Huang
2018-02-28

Leverage the power of Redis 4.x to develop, optimize and administer your Redis solutions with ease Key Features Build, deploy and administer high performance and scalable applications in Redis Covers a range of important tasks -

including development and administration of Redis A practical guide that takes your understanding of Redis to the next level Book Description Redis is considered the world's most popular key-value store database. Its versatility and the wide variety of use cases it enables have made it a popular choice of database for many enterprises. Based on the latest version of Redis, this book provides both step-by-step recipes and relevant the background information required to utilize its features to the fullest. It covers everything from a basic understanding of Redis data types to advanced aspects of Redis high availability, clustering, administration, and troubleshooting. This book will be your great companion to master all aspects of Redis. The book starts off by installing and configuring Redis for you to get started with ease. Moving on, all the data types and features of Redis are introduced in detail. Next, you will learn how to develop applications with Redis in Java, Python, and the Spring Boot web framework. You will also learn replication tasks, which will help you to troubleshoot replication issues. Furthermore, you will learn the steps that need to be undertaken to ensure high availability on your cluster and during production deployment. Toward the end of the book, you will learn the topmost tasks that will help you to troubleshoot your ecosystem efficiently, along with extending Redis by using different modules. What you will learn Install and configure your Redis instance Explore various data types and commands in Redis Build client-side applications as well as a Big Data framework with Redis Manage data replication and persistence in Redis Implement high availability and data sharding in Redis Extend Redis with Redis Module Benchmark, debug, fine-tune and troubleshoot various issues in Redis Who this book is for This book is for database administrators, developers and architects who want to tackle the common and not so common problems associated with the different development and administration-related tasks in Redis. A fundamental understanding of Redis is expected to get the best out of this book.

Accessibility Handbook - Katie Cunningham 2012-08-27

Get practical guidelines for making your website accessible to people with disabilities. With this handbook, you'll learn how to design or develop

a site that conforms to Section 508 of the US Rehabilitation Act—and in the process you'll discover how to provide a better user experience for everyone. The Accessibility Handbook introduces you to several audiences that have difficulty using today's complex websites, including people with blindness, hearing loss, physical disabilities, and cognitive disorders. Learn how to support assistive technologies, and understand which fonts, colors, page layouts, and other design elements work best—without having to exclude advanced functions, hire outside help, or significantly increase overhead. Develop solutions that accommodate: Complete blindness. Create a logical document flow to support screen readers Low vision and color blindness. Optimize images and color schemes, and ensure your site enlarges gracefully Hearing impairment. Provide video captions and visual alerts for interactive features Physical disabilities. Make forms, popups, and navigation easier to use Cognitive disorders. Adapt fonts and text styles for dyslexic users, and design consistent, well-organized pages for people with ADHD

Getting Started with Storm - Jonathan Leibiusky 2012-08-31

Even as big data is turning the world upside down, the next phase of the revolution is already taking shape: real-time data analysis. This hands-on guide introduces you to Storm, a distributed, JVM-based system for processing streaming data. Through simple tutorials, sample Java code, and a complete real-world scenario, you'll learn how to build fast, fault-tolerant solutions that process results as soon as the data arrives. Discover how easy it is to set up Storm clusters for solving various problems, including continuous data computation, distributed remote procedure calls, and data stream processing. Learn how to program Storm components: spouts for data input and bolts for data transformation Discover how data is exchanged between spouts and bolts in a Storm topology Make spouts fault-tolerant with several commonly used design strategies Explore bolts—their life cycle, strategies for design, and ways to implement them Scale your solution by defining each component's level of parallelism Study a real-time web analytics system built with Node.js, a Redis server, and a Storm topology Write spouts and bolts with non-JVM languages such as Python, Ruby, and

Javascript

Spring Boot in Action - Craig Walls 2015-12-16

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

Spring Data - Mark Pollack 2012-10-24

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

Learning Redis - Vinoo Das 2015-06-26

This book will introduce Redis and help you understand its various facets. Starting with an introduction to NoSQL, you will learn how to install Redis and how to classify and work with data structures. By working with real world scenarios pertaining to using Redis, you will discover sharding and indexing techniques, along with how to improve scalability and performance through persistent strategies and data migration techniques. With the help of multiple examples, you will learn to design web and business applications. You will also learn how to configure Redis for setting up clusters and tuning it for performance. At the end of this book, you will find essential tips on backup and recovery strategies for the Redis environment.

Vert.x in Action - Julien Ponge 2020-12-01

Vert.x in Action teaches you how to build production-quality reactive applications in Java. This book covers core Vert.x concepts, as well as the fundamentals of asynchronous and reactive programming. Learn to develop microservices by

using Vert.x tools for database communications, persistent messaging, and test app resiliency. The patterns and techniques included here transfer to reactive technologies and frameworks beyond Vert.x. Summary As enterprise applications become larger and more distributed, new architectural approaches like reactive designs, microservices, and event streams are required knowledge. The Vert.x framework provides a mature, rock-solid toolkit for building reactive applications using Java, Kotlin, or Scala. Vert.x in Action teaches you to build responsive, resilient, and scalable JVM applications with Vert.x using well-established reactive design patterns. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Vert.x is a collection of libraries for the Java virtual machine that simplify event-based and asynchronous programming. Vert.x applications handle tedious tasks like asynchronous communication, concurrent work, message and data persistence, plus they're easy to scale, modify, and maintain. Backed by the Eclipse Foundation and used by Red Hat and others, this toolkit supports code in a variety of languages. About the book Vert.x in Action teaches you how to build production-quality reactive applications in Java. This book covers core Vert.x concepts, as well as the fundamentals of asynchronous and reactive programming. Learn to develop microservices by using Vert.x tools for database communications, persistent messaging, and test app resiliency. The patterns and techniques included here transfer to reactive technologies and frameworks beyond Vert.x. What's inside Building reactive services Responding to external service failures Horizontal scaling Vert.x toolkit architecture and Vert.x testing Deploying with Docker and Kubernetes About the reader For intermediate Java web developers. About the author Julien Ponge is a principal software engineer at Red Hat, working on the Eclipse Vert.x project. Table of Contents PART 1 - FUNDAMENTALS OF ASYNCHRONOUS PROGRAMMING WITH VERT.X 1 Vert.x, asynchronous programming, and reactive systems 2 Verticles: The basic processing units of Vert.x 3 Event bus: The backbone of a Vert.x application 4 Asynchronous data and event streams 5 Beyond callbacks 6 Beyond the event

bus PART 2 - DEVELOPING REACTIVE SERVICES WITH VERT.X 7 Designing a reactive application 8 The web stack 9 Messaging and event streaming with Vert.x 10 Persistent state management with databases 11 End-to-end real-time reactive event processing 12 Toward responsiveness with load and chaos testing 13 Final notes: Container-native Vert.x MongoDB in Action - Kyle Banker 2016-03-29 Summary MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology This document-oriented database was built for high availability, supports rich, dynamic schemas, and lets you easily distribute data across multiple servers. MongoDB 3.0 is flexible, scalable, and very fast, even with big data loads. About the Book MongoDB in Action, Second Edition is a completely revised and updated version. It introduces MongoDB 3.0 and the document-oriented database model. This perfectly paced book gives you both the big picture you'll need as a developer and enough low-level detail to satisfy system engineers. Lots of examples will help you develop confidence in the crucial area of data modeling. You'll also love the deep explanations of each feature, including replication, auto-sharding, and deployment. What's Inside Indexes, queries, and standard DB operations Aggregation and text searching Map-reduce for custom aggregations and reporting Deploying for scale and high availability Updated for Mongo 3.0 About the Reader Written for developers. No previous MongoDB or NoSQL experience is assumed. About the Authors After working at MongoDB, Kyle Banker is now at a startup. Peter Bakkum is a developer with MongoDB expertise. Shaun Verch has worked on the core server team at MongoDB. A Genentech engineer, Doug Garrett is one of the winners of the MongoDB Innovation Award for Analytics. A software architect, Tim Hawkins has led search engineering at Yahoo Europe. Technical Contributor: Wouter Thielen. Technical Editor: Mihalis Tsoukalos. Table of Contents PART 1

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

Redis Cookbook - Tiago Macedo 2011-07-25

Two years since its initial release, Redis already has an impressive list of adopters, including Engine Yard, GitHub, Craigslist, and Digg. This open source data structure server is built for speed and flexibility, making it ideal for many applications. If you're using Redis, or considering it, this concise cookbook provides recipes for a variety of issues you're likely to face. Each recipe solves a specific problem, and provides an in-depth discussion of how the solution works. You'll discover that Redis, while simple in nature, offers extensive functionality for manipulating and storing data. Learn when it makes sense to use Redis Explore several methods for installing Redis Connect to Redis in a number of ways, ranging from the command line to popular languages such as Python and Ruby Solve a range of needs, from linked datasets to analytics Handle backups, sharding, datasets larger than available memory, and many other tasks

Mastering Redis - Jeremy Nelson 2016-05-31

Take your knowledge of Redis to the next level to build enthralling applications with ease About This Book Detailed explanation on Data structure server with powerful strings, lists, sets, sorted-sets, and hashes Learn to Scale your data with Redis Cluster's distributed setup This is a fast paced practical guide full of screenshots and real work examples to help you get to grips with Redis in no time. Who This Book Is For If you are a software developer with some experience with Redis and would now like to elevate your Redis knowledge and skills even further, then this book is for you. What You Will Learn Choose the right Redis data structure for your problem Understand Redis event-loop and implement your own custom C commands Solve complex workflows with Redis server-side scripting with Lua Configure your Redis instance for optimal

memory management Scale your data in a distributed manner with Redis Cluster Improve the stability of your Redis solution using Redis Sentinel Complement your existing database and NoSQL environment with Redis Exploit a wide range of features provided by Redis to become a DevOps expert. In Detail Redis is the most popular, open-source, key value data structure server that provides a wide range of capabilities on which multiple platforms can be built. Its fast and flexible data structures give your existing applications an edge in the development environment. This book is a practical guide which aims to help you deep dive into the world of Redis data structure to exploit its excellent features. We start our journey by understanding the need of Redis in brief, followed by an explanation of Advanced key management. Next, you will learn about design patterns, best practices for using Redis in DevOps environment and Docker containerization paradigm in detail. After this, you will understand the concept of scaling with Redis cluster and Redis Sentinel , followed by a through explanation of incorporating Redis with NoSQL technologies such as Elasticsearch and MongoDB. At the end of this section, you will be able to develop competent applications using these technologies. You will then explore the message queuing and task management features of Redis and will be able to implement them in your applications. Finally, you will learn how Redis can be used to build real-time data analytic dashboards, for different disparate data streams. Style and approach This is a hands on guide full of easy-to-follow examples, that illustrate important concepts and techniques to solve complex problems with Redis.

GraphQL in Action - Samer Buna 2021-03-09

GraphQL in Action gives you the tools to get comfortable with the GraphQL language, build and optimize a data API service, and use it in a front-end client application. Summary Reduce bandwidth demands on your APIs by getting only the results you need—all in a single request! The GraphQL query language simplifies interactions with web servers, enabling smarter API queries that can hugely improve the efficiency of data requests. In GraphQL in Action, you'll learn how to bring those benefits to your own APIs, giving your clients the power to ask for exactly what

they need from your server, no more, no less. Practical and example-driven, this book teaches everything you need to get started with GraphQL—from design principles and syntax right through to performance optimization. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology GraphQL APIs are fast, efficient, and easy to maintain. They reduce app latency and server cost while boosting developer productivity. This powerful query layer offers precise control over API requests and returns, making apps faster and less prone to error. About the book GraphQL in Action gives you the tools to get comfortable with the GraphQL language, build and optimize a data API service, and use it in a front-end client application. By working through set up, security, and error handling you'll learn to create a complete GraphQL server. You'll also unlock easy ways to incorporate GraphQL into your existing codebase so you can build simple, scalable data APIs. What's inside Define a GraphQL schema for relational and document databases Implement GraphQL types using both the schema language and object constructor methods Optimize GraphQL resolvers with data caching and batching Design GraphQL fragments that match UI components' data requirements Consume GraphQL API queries, mutations, and subscriptions with and without a GraphQL client library About the reader For web developers familiar with client-server applications. About the author Samer Buna has over 20 years of experience in software development including front-ends, back-ends, API design, and scalability. Table of Contents PART 1- EXPLORING GRAPHQL 1 Introduction to GraphQL 2 Exploring GraphQL APIs 3 Customizing and organizing GraphQL operations PART 2 - BUILDING GRAPHQL APIs 4 Designing a GraphQL schema 5 Implementing schema resolvers 6 Working with database models and relations 7 Optimizing data fetching 8 Implementing mutations PART 3 - USING GRAPHQL APIs 9 Using GraphQL APIs without a client library 10 Using GraphQL APIs with Apollo client *Seven Databases in Seven Weeks* - Luc Perkins 2018-04-05 Data is getting bigger and more complex by the day, and so are your choices in handling it.

Explore some of the most cutting-edge databases available - from a traditional relational database to newer NoSQL approaches - and make informed decisions about challenging data storage problems. This is the only comprehensive guide to the world of NoSQL databases, with in-depth practical and conceptual introductions to seven different technologies: Redis, Neo4J, CouchDB, MongoDB, HBase, Postgres, and DynamoDB. This second edition includes a new chapter on DynamoDB and updated content for each chapter. While relational databases such as MySQL remain as relevant as ever, the alternative, NoSQL paradigm has opened up new horizons in performance and scalability and changed the way we approach data-centric problems. This book presents the essential concepts behind each database alongside hands-on examples that make each technology come alive. With each database, tackle a real-world problem that highlights the concepts and features that make it shine. Along the way, explore five database models - relational, key/value, columnar, document, and graph - from the perspective of challenges faced by real applications. Learn how MongoDB and CouchDB are strikingly different, make your applications faster with Redis and more connected with Neo4J, build a cluster of HBase servers using cloud services such as Amazon's Elastic MapReduce, and more. This new edition brings a brand new chapter on DynamoDB, updated code samples and exercises, and a more up-to-date account of each database's feature set. Whether you're a programmer building the next big thing, a data scientist seeking solutions to thorny problems, or a technology enthusiast venturing into new territory, you will find something to inspire you in this book. What You Need: You'll need a *nix shell (Mac OS or Linux preferred, Windows users will need Cygwin), Java 6 (or greater), and Ruby 1.8.7 (or greater). Each chapter will list the downloads required for that database.

Mastering GitLab 12 - Joost Evertse
2019-08-02

An expert guide to helping you use DevOps techniques with the latest GitLab version to optimize and manage your software workflow Key FeaturesDelve into GitLab's architecture, and install and configure it to fit your

environment Learn about the underlying principles of Agile software development and DevOps Explore Gitlab's features to manage enterprise cloud-native applications and services Book Description GitLab is an open source repository management and version control toolkit with functions for enterprises and personal software projects. It offers configurability options, extensions, and APIs that make it an ideal tool for enterprises to manage the software development life cycle. This book begins by explaining GitLab options and the components of the GitLab architecture. You will learn how to install and set up GitLab on-premises and in the cloud, along with understanding how to migrate code bases from different systems, such as GitHub, Concurrent Versions System, Team Foundation Version Control, and Subversion. Later chapters will help you implement DevOps culture by introducing the workflow management tools in GitLab and continuous integration/continuous deployment (CI/CD). In addition to this, the book will guide you through installing GitLab on a range of cloud platforms, monitoring with Prometheus, and deploying an environment with GitLab. You'll also focus on the GitLab CI component to assist you with creating development pipelines and jobs, along with helping you set up GitLab runners for your own project. Finally, you will be able to choose a high availability setup that fits your needs and helps you monitor and act on results obtained after testing. By the end of this book, you will have gained the expertise you need to use GitLab features effectively, and be able to integrate all phases in the development process. What you will learn Install GitLab on premises and in the cloud using a variety of configurations Conduct data migration from the SVN, TFS, CVS, and GitHub platforms to GitLab Use GitLab runners to develop different types of configurations in software development Plan and perform CI/CD by using GitLab features Monitor and secure your software architecture using Prometheus and Grafana Implement DevOps culture by introducing workflow management tools in GitLab Who this book is for If you are a software developer, DevOps professional, or any developer who wants to master GitLab for productive repository management in your day-to-day tasks, this book

is for you. Basic understanding of the software development workflow is assumed.

Willpower - Roy F. Baumeister 2011-09-01 One of the world's most esteemed and influential psychologists, Roy F. Baumeister, teams with New York Times science writer John Tierney to reveal the secrets of self-control and how to master it. "Deep and provocative analysis of people's battle with temptation and masterful insights into understanding willpower: why we have it, why we don't, and how to build it. A terrific read." —Ravi Dhar, Yale School of Management, Director of Center for Customer Insights Pioneering research psychologist Roy F. Baumeister collaborates with New York Times science writer John Tierney to revolutionize our understanding of the most coveted human virtue: self-control. Drawing on cutting-edge research and the wisdom of real-life experts, Willpower shares lessons on how to focus our strength, resist temptation, and redirect our lives. It shows readers how to be realistic when setting goals, monitor their progress, and how to keep faith when they falter. By blending practical wisdom with the best of recent research science, Willpower makes it clear that whatever we seek—from happiness to good health to financial security—we won't reach our goals without first learning to harness self-control.

Web Caching - Duane Wessels 2001 On the World Wide Web, speed and efficiency are vital. Users have little patience for slow web pages, while network administrators want to make the most of their available bandwidth. A properly designed web cache reduces network traffic and improves access times to popular web sites—a boon to network administrators and web users alike. Web Caching hands you all the technical information you need to design, deploy, and operate an effective web caching service. It starts with the basics of how web caching works, from the HTTP headers that govern cachability to cache validation and replacement algorithms. Topics covered in this book include: Designing an effective cache solution Configuring web browsers to use a cache Setting up a collection of caches that can talk to each other Configuring an interception cache or proxy Monitoring and fine-tuning the performance of a cache Configuring web servers to cooperate with web caches Benchmarking cache products The

book also covers the important political aspects of web caching, including privacy, intellectual property, and security issues. Internet service providers, large corporations, or educational institutions--in short, any network that provides connectivity to a wide variety of users--can reap enormous benefit from running a well-tuned web caching service. Web Caching shows you how to do it right.

Play for Scala - Peter Hilton 2013-10-03

Summary Play for Scala shows you how to build Scala-based web applications using the Play 2 framework. This book starts by introducing Play through a comprehensive overview example. Then, you'll look at each facet of a typical Play application both by exploring simple code snippets and by adding to a larger running example. Along the way, you'll deepen your knowledge of Scala as a programming language and work with tools like Akka. About this Book Play is a Scala web framework with built-in advantages: Scala's strong type system helps deliver bug-free code, and the Akka framework helps achieve hassle-free concurrency and peak performance. Play builds on the web's stateless nature for excellent scalability, and because it is event-based and nonblocking, you'll find it to be great for near real-time applications. Play for Scala teaches you to build Scala-based web applications using Play 2. It gets you going with a comprehensive overview example. It then explores each facet of a typical Play application by walking through sample code snippets and adding features to a running example. Along the way, you'll deepen your knowledge of Scala and learn to work with tools like Akka. Written for readers familiar with Scala and web-based application architectures. No knowledge of Play is assumed. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside Intro to Play 2 Play's MVC structure Mastering Scala templates and forms Persisting data and using web services Using Play's advanced features About the Authors Peter Hilton, Erik Bakker, and Francisco Canedo, are engineers at Lunatech, a consultancy with Scala and Play expertise. They are contributors to the Play framework. Table of Contents PART 1: GETTING STARTED Introduction to Play Your first Play application PART 2: CORE FUNCTIONALITY Deconstructing Play application

architecture Defining the application's HTTP interface Storing data—the persistence layer Building a user interface with view templates Validating and processing input with the forms API PART 3: ADVANCED CONCEPTS Building a single-page JavaScript application with JSON Play and more Web services, iteratees, and WebSockets

Big Data - Nathan Marz 2015

Summary Big Data teaches you to build big data systems using an architecture that takes advantage of clustered hardware along with new tools designed specifically to capture and analyze web-scale data. It describes a scalable, easy-to-understand approach to big data systems that can be built and run by a small team. Following a realistic example, this book guides readers through the theory of big data systems, how to implement them in practice, and how to deploy and operate them once they're built. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book Web-scale applications like social networks, real-time analytics, or e-commerce sites deal with a lot of data, whose volume and velocity exceed the limits of traditional database systems. These applications require architectures built around clusters of machines to store and process data of any size, or speed. Fortunately, scale and simplicity are not mutually exclusive. Big Data teaches you to build big data systems using an architecture designed specifically to capture and analyze web-scale data. This book presents the Lambda Architecture, a scalable, easy-to-understand approach that can be built and run by a small team. You'll explore the theory of big data systems and how to implement them in practice. In addition to discovering a general framework for processing big data, you'll learn specific technologies like Hadoop, Storm, and NoSQL databases. This book requires no previous exposure to large-scale data analysis or NoSQL tools. Familiarity with traditional databases is helpful. What's Inside Introduction to big data systems Real-time processing of web-scale data Tools like Hadoop, Cassandra, and Storm Extensions to traditional database skills About the Authors Nathan Marz is the creator of Apache Storm and the originator of the Lambda Architecture for big data systems. James Warren

is an analytics architect with a background in machine learning and scientific computing. Table of Contents A new paradigm for Big Data PART 1 BATCH LAYER Data model for Big Data Data model for Big Data: Illustration Data storage on the batch layer Data storage on the batch layer: Illustration Batch layer Batch layer: Illustration An example batch layer: Architecture and algorithms An example batch layer: Implementation PART 2 SERVING LAYER Serving layer Serving layer: Illustration PART 3 SPEED LAYER Realtime views Realtime views: Illustration Queuing and stream processing Queuing and stream processing: Illustration Micro-batch stream processing Micro-batch stream processing: Illustration Lambda Architecture in depth

Cassandra: The Definitive Guide - Jeff Carpenter 2016-06-29

Imagine what you could do if scalability wasn't a problem. With this hands-on guide, you'll learn how the Cassandra database management system handles hundreds of terabytes of data while remaining highly available across multiple data centers. This expanded second edition—updated for Cassandra 3.0—provides the technical details and practical examples you need to put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's non-relational design, with special attention to data modeling. If you're a developer, DBA, or application architect looking to solve a database scaling issue or future-proof your application, this guide helps you harness Cassandra's speed and flexibility. Understand Cassandra's distributed and decentralized structure Use the Cassandra Query Language (CQL) and cqlsh—the CQL shell Create a working data model and compare it with an equivalent relational model Develop sample applications using client drivers for languages including Java, Python, and Node.js Explore cluster topology and learn how nodes exchange data Maintain a high level of performance in your cluster Deploy Cassandra on site, in the Cloud, or with Docker Integrate Cassandra with Spark, Hadoop, Elasticsearch, Solr, and Lucene

Squid: The Definitive Guide - Duane Wessels 2004-01-22

Squid is the most popular Web caching software in use today, and it works on a variety of

platforms including Linux, FreeBSD, and Windows. Squid improves network performance by reducing the amount of bandwidth used when surfing the Web. It makes web pages load faster and can even reduce the load on your web server. By caching and reusing popular web content, Squid allows you to get by with smaller network connections. It also protects the host on your internal network by acting as a firewall and proxying your internal web traffic. You can use Squid to collect statistics about the traffic on your network, prevent users from visiting inappropriate web sites at work or school, ensure that only authorized users can surf the Internet, and enhance your privacy by filtering sensitive information from web requests. Companies, schools, libraries, and organizations that use web-caching proxies can look forward to a multitude of benefits. Written by Duane Wessels, the creator of Squid, Squid: The Definitive Guide will help you configure and tune Squid for your particular situation. Newcomers to Squid will learn how to download, compile, and install code. Seasoned users of Squid will be interested in the later chapters, which tackle advanced topics such as high-performance storage options, rewriting requests, HTTP server acceleration, monitoring, debugging, and troubleshooting Squid. Topics covered include: Compiling and installing Squid Running Squid Using Squid's sophisticated access controls Tuning disk storage for optimal performance Configuring your operating system for HTTP interception Forwarding Requests to other web caches Using redirectors to rewrite user requests Monitoring Squid with the cache manager and SNMP Using Squid to accelerate and protect HTTP servers Managing bandwidth consumption with Delay Pools

Java Enterprise Best Practices - O'Reilly Java Authors 2002-12-18

Java developers typically go through four "stages" in mastering Java. In the first stage, they learn the language itself. In the second stage, they study the APIs. In the third stage, they become proficient in the environment. It is in the fourth stage --"the expert stage"-- where things really get interesting, and Java Enterprise Best Practices is the tangible compendium of experience that developers need to breeze through this fourth and final stage of Enterprise Java mastery. Crammed with tips and tricks, Java

Enterprise Best Practices distills years of solid experience from eleven experts in the J2EE environment into a practical, to-the-point guide to J2EE. Java Enterprise Best Practices gives developers the unvarnished, expert-tested advice that the man pages don't provide--what areas of the APIs should be used frequently (and which are better avoided); elegant solutions to problems you face that other developers have already discovered; what things you should always do, what things you should consider doing, and what things you should never do--even if the documentation says it's ok. Until Java Enterprise Best Practices, Java developers in the fourth stage of mastery relied on the advice of a loose-knit community of fellow developers, time-consuming online searches for examples or suggestions for the immediate problem they faced, and tedious trial-and-error. But Java has grown to include a huge number of APIs, classes, and methods. Now it is simply too large for even the most intrepid developer to know it all. The need for a written compendium of J2EE Best Practices has never been greater. Java Enterprise Best Practices focuses on the Java 2 Enterprise Edition (J2EE) APIs. The J2EE APIs include such alphabet soup acronyms as EJB, JDBC, RMI, XML, and JMX.

Seven NoSQL Databases in a Week - Xun (Brian) Wu 2018-03-29

A beginner's guide to get you up and running with Cassandra, DynamoDB, HBase, InfluxDB, MongoDB, Neo4j, and Redis. Key Features: Covers the basics of 7 NoSQL databases and how they are used in the enterprises. Quick introduction to MongoDB, DynamoDB, Redis, Cassandra, Neo4j, InfluxDB, and HBase. Includes effective techniques for database querying and management. Book Description: This is the golden age of open source NoSQL databases. With enterprises having to work with large amounts of

unstructured data and moving away from expensive monolithic architecture, the adoption of NoSQL databases is rapidly increasing. Being familiar with the popular NoSQL databases and knowing how to use them is a must for budding DBAs and developers. This book introduces you to the different types of NoSQL databases and gets you started with seven of the most popular NoSQL databases used by enterprises today. We start off with a brief overview of what NoSQL databases are, followed by an explanation of why and when to use them. The book then covers the seven most popular databases in each of these categories: MongoDB, Amazon DynamoDB, Redis, HBase, Cassandra, InfluxDB, and Neo4j. The book doesn't go into too much detail about each database but teaches you enough to get started with them. By the end of this book, you will have a thorough understanding of the different NoSQL databases and their functionalities, empowering you to select and use the right database according to your needs. What you will learn: Understand how MongoDB provides high-performance, high-availability, and automatic scaling. Interact with your Neo4j instances via database queries, Python scripts, and Java application code. Get familiar with common querying and programming methods to interact with Redis. Study the different types of problems Cassandra can solve. Work with HBase components to support common operations such as creating tables and reading/writing data. Discover data models and work with CRUD operations using DynamoDB. Discover what makes InfluxDB a great choice for working with time-series data. Who this book is for: If you are a budding DBA or a developer who wants to get started with the fundamentals of NoSQL databases, this book is for you. Relational DBAs who want to get insights into the various offerings of popular NoSQL databases will also find this book to be very useful.