

Pro Java Clustering And Scalability Building Real

Foundations of Scalable Systems
 Reactive Programming with Scala and Akka
 The Art of Multiprocessor Programming, Revised Reprint
 InfoWorld
 Redis Essentials
 Pro Java Microservices with Quarkus and Kubernetes
 PRO JSK SITE D,
 Mahout in Action
 Annual Review of Scalable Computing
 Handbook of Cloud Computing
 Java Performance and Scalability: Server-side programming techniques
 Effective Kafka
 Pro Oracle Collaboration Suite 10g
 Data-Intensive Text Processing with MapReduce
 Annual Review of Scalable Computing
 Getting Started with Hazelcast - Second Edition
 WebSocket
 Mastering Microservices with Java 9
 Practical Node.js
 Pro Java EE 5 Performance Management and Optimization
 Building Scalable Applications with Erlang
 Pro (IBM) WebSphere Application Server 7 Internals
 The Art of Scalability
 Annual Review of Scalable Computing
 Infinispan Data Grid Platform Definitive Guide
 Pro Spring Integration
 Mining of Massive Datasets
 Mastering Enterprise JavaBeans
 The Definitive Guide to Terracotta
 Pro JavaFX 9
 WebSocket Essentials - Building Apps with HTML5 WebSockets
 Distributed and Cloud Computing
 Java Performance and Scalability
 Designing Data-Intensive Applications
 Pro Java Clustering and Scalability
 Scalable Data Architecture with Java
 Pro Spring 5
 Dr. Dobb's Journal
 Building Microservices Applications on Microsoft Azure
 Oracle Coherence 3.5

Pro Java Clustering And Scalability Building Real

Downloaded from data.avac.org by guest

PRECIOUS CARLEE

Foundations of Scalable Systems

Packt Publishing Ltd
 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.

Reactive Programming with Scala and Akka

World Scientific
 Now in its second edition, this book focuses on practical algorithms for mining data from even the largest datasets.

The Art of Multiprocessor Programming, Revised Reprint

Packt Publishing Ltd
 Includes more than 30 percent revised material and five new chapters, covering the new 2.1 features such as EJB Timer Service and JMS as well as the latest open source Java solutions The book was developed as part of TheServerSide.com online EJB community, ensuring a built-in audience Demonstrates how to build an EJB system, program with EJB, adopt best practices, and harness advanced EJB concepts and techniques, including transactions, persistence, clustering, integration, and performance optimization Offers practical guidance on when not to use EJB and how to use simpler, less costly open source technologies in place of or in conjunction with EJB

InfoWorld Apress

This book is a great introduction for Java developers, software architects, or DevOps looking to enable scalable and agile data within their applications. Providing in-memory object storage, cluster-wide state and messaging, or even scalable task execution, Hazelcast helps solve a number of issues that have troubled technologists for years.

Redis Essentials Apress

In many systems, scalability becomes the primary driver as the user base grows. Attractive features and high utility breed success, which brings more requests to handle and more data to manage. But organizations reach a tipping point when design decisions that made sense under light loads suddenly become technical debt. This practical book covers design approaches and technologies that make it possible to scale an application quickly and cost-effectively. Author Ian Gorton takes software architects and developers through the foundational principles of distributed systems. You'll explore the essential ingredients of scalable solutions, including replication, state management, load balancing, and caching. Specific chapters focus on the implications of scalability for databases, microservices, and event-based streaming systems. You will focus on: Foundations of scalable systems: Learn basic design principles of scalability, its costs, and architectural tradeoffs Designing scalable services: Dive into service design, caching, asynchronous messaging, serverless processing, and microservices Designing scalable data systems: Learn data system fundamentals, NoSQL databases, and eventual consistency versus strong consistency Designing scalable streaming systems: Explore stream processing systems and scalable event-driven processing

Pro Java Microservices with Quarkus and Kubernetes

Apress
 Pro (IBM) WebSphere Application Server 7 Internals covers the internal architecture and implementation of the WebSphere Application Server (WAS) version 7 product set and how other IBM products extend it. It presents information to enable administrators, developers, and architects to learn about the aspects of WAS that apply to them: Administrators will come to understand how the WAS7 environment functions to best optimize it for their environment, and what to do when things go wrong. Developers will learn to extend the functionality in the base WAS product. Architects will see how the WAS product underpins the IBM offerings to fit in an enterprise.

PRO JSK SITE D, World Scientific

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures

Mahout in Action

Createspace Independent Pub
 Pro Spring Integration is an authoritative book from the experts that guides you through the vast world of enterprise application integration (EAI) and application of the Spring Integration framework towards solving integration problems. The book is: An introduction to the concepts of enterprise application integration A reference on building event-driven applications using Spring Integration A guide to solving common integration problems using Spring Integration What makes this book unique is its coverage of contemporary technologies and real-world information, with a focus on common problems that users are likely to confront. This book zeroes in on extending the Spring Integration framework to meet your custom integration demands. As Spring Integration is an extension of the Spring programming model, it builds on the Spring Framework's existing support for enterprise integration. This book will take you through all aspects of this relationship and show you

how to get the most out of your Spring applications, where integration is a consideration. It discusses simple messaging within Spring-based applications and integration with external systems via simple adapters. Those adapters provide a higher-level of abstraction over Spring's support for remoting, messaging, and scheduling, all of which receives coverage in this book.

Annual Review of Scalable Computing Apress

This practical guide is intended for those who want to learn how to build extremely scalable applications. This book is easy to read and is aimed at Java enterprise developers with a solid knowledge of Java. However, no previous coding experience with Infinispan is required.

Handbook of Cloud Computing Packt Publishing Ltd

Written in Henry Liu's clear, concise style, *Java Performance and Scalability* gets right to the point. With clearly explained concepts, most pertinent theories, precise step-by-step procedures, and large volume of illustrative charts and tables with highly reliable data supporting behind, you gain quickly the necessary knowledge and skills for being able to cope with Java application performance and scalability issues without having to resort to more experienced professionals or expensive external consultants. Specifically, it helps you learn the following knowledge and skills that are essential for you to become more effective in contributing to the success of your organization: * What you need to know at minimum about the architecture of modern hardware so that you can make smart decisions on when you should pour your time on your application and when you can just throw in more advanced hardware to get by. * What you need to know about garbage collection theories in general and how they are implemented with widely used Java Virtual Machines like HotSpot JVMs. * Precise methodologies, procedures, and programs that you can start to use immediately to help you profile and tune your Java applications. * How you can design and build performance and scalability into your product proactively without having to face tough retrofitting decisions or even torrents of customer escalations later on. * Optimizing and tuning Java performance and scalability on Linux with comparison between Linux and Windows. * CPU frequency scaling benefits and side effects with Intel's Turbo Boost Technology on Linux and Windows. In addition, the book contains interesting data for your reference, associated with oops compression, CMS garbage collection tuning, DoEscapeAnalysis, G1 versus CMS comparison, Linux versus Windows, CPU frequency scaling benefits and side effects with Intel's Turbo Boost Technology on Linux and Windows, etc., all based on full scale, rigorous performance and scalability tests with real products.

Java Performance and Scalability: Server-side programming techniques Elsevier

This book is for web developers who want to learn and implement WebSocket to create interesting apps for modern browsers, leveraging the capabilities of HTML5 with WebSockets.

Effective Kafka "O'Reilly Media, Inc."

Revised and updated with improvements conceived in parallel programming courses, *The Art of Multiprocessor Programming* is an authoritative guide to multicore programming. It introduces a higher level set of software development skills than that needed for efficient single-core programming. This book provides comprehensive coverage of the new principles, algorithms, and tools necessary for effective multiprocessor programming. Students and professionals alike will benefit from thorough coverage of key multiprocessor programming issues. This revised edition incorporates much-demanded updates throughout the book, based on feedback and corrections reported from classrooms since 2008. Learn the fundamentals of programming multiple threads accessing shared memory. Explore mainstream concurrent data structures and the key elements of their design, as well as synchronization techniques from simple locks to transactional memory systems. Visit the companion site and download source code, example Java programs, and materials to support and enhance the learning experience.

Pro Oracle Collaboration Suite 10g Packt Publishing Ltd

This book contains four review articles in the area of scalable computing. Two of the articles discuss methods and tools for the parallel solution of irregular problems, which have been satisfactorily worked out in heterogeneous systems. One surveys the technology and applications of multimedia server clusters, which are playing an increasing role in the current networked environment. An additional article discusses SilkRoad, which adds distributed shared memory capabilities to the Cilk parallel programming system. Once again, the book represents a new set of steps forward in parallel systems. Graduate students, academics and researchers in supercomputing and computer engineering.

Data-Intensive Text Processing with MapReduce Addison-Wesley Professional

Build and design microservices using Java and the Red Hat Quarkus Framework. This book will help you quickly get started with the features and concerns of a microservices architecture. It will introduce Docker and Kubernetes to help you deploy your microservices. You will be guided on how to install the appropriate tools to work properly. For those who are new to enterprise development using Quarkus, you will be introduced to its core principles and main features through a deep step-by-step tutorial. For experts, this book offers some recipes that illustrate how to split monoliths and implement microservices and deploy them as containers to Kubernetes. By the end of reading this book, you will have practical hands-on experience of building microservices using Quarkus and you will master deploying them to Kubernetes. You will: Work with Quarkus and GraalVM Split a monolith using the domain-driven design approach Implement the cloud and microservices patterns Rethink the deployment process Introduce containerization, Docker, and Kubernetes to your toolkit Boost microservices efficiency and performance with Azure Play with Quarkus and distributed application runtimes.

Annual Review of Scalable Computing Packt Publishing Ltd

Implement microservices starting with their architecture and moving on to their deployment, manageability, security, and monitoring. This book focuses on the key scenarios where microservices architecture is preferred over a monolithic architecture. *Building Microservices Applications on Microsoft Azure* begins with a survey of microservices architecture compared to monolithic architecture and covers microservices implementation in detail. You'll see the key scenarios where microservices architecture is preferred over a monolithic approach. From there, you will explore the critical components and various deployment options of microservices on platforms such as Microsoft Azure (public cloud) and Azure Stack (hybrid cloud). This includes in-depth coverage of developing, deploying, and monitoring microservices on containers and orchestrating with Azure Service Fabric and Azure Kubernetes Cluster (AKS). This book includes practical experience from large-scale enterprise deployments, therefore it can be a quick reference for solution architects and developers to understand the critical factors while designing a microservices application. *What You Will Learn* Explore the use cases of microservices and monolithic architecture Discover the architecture patterns to build scalable, agile, and secure microservices applications Develop and deploy microservices using Azure Service Fabric and Azure Kubernetes Service Secure microservices using the gateway pattern See the deployment options for Microservices on Azure Stack Implement database patterns to handle the complexities introduced by microservices Who This Book Is For Architects and consultants who work on Microsoft Azure and manage large-scale deployments.

Getting Started with Hazelcast - Second Edition Springer Nature

Practical Node.js is your step-by-step guide to learning how to build a wide range of scalable real-world web applications using a professional development toolkit. Node.js is an innovative and highly

efficient platform for creating web services. But Node.js doesn't live in a vacuum! In a modern web development, many different components need to be put together — routing, database driver, ORM, session management, OAuth, HTML template engine, CSS compiler and many more. If you already know the basics of Node.js, now is the time to discover how to bring it to production level by leveraging its vast ecosystem of packages. As a web developer, you'll work with a varied collection of standards and frameworks - *Practical Node.js* shows you how all those pieces fit together.

Practical Node.js takes you from installing all the necessary modules to writing full-stack web applications by harnessing the power of the Express.js and Hapi frameworks, the MongoDB database with Mongoskin and Mongoose, Jade and Handlebars template engines, Stylus and LESS CSS languages, OAuth and Everyauth libraries, and the Socket.IO and Derby libraries, and everything in between. The book also covers how to deploy to Heroku and AWS, daemonize apps, and write REST APIs. You'll build full-stack real-world Node.js apps from scratch, and also discover how to write your own Node.js modules and publish them on NPM. You already know what Node.js is; now learn what you can do with it and how far you can take it!

WebSocket Springer Science & Business Media

Terracotta is a High Availability (HA) nth degree scaling and clustering engine for traditional J2EE and Java EE 5 applications (using Seam or other) as well as Spring-based enterprise applications. Written and officially authorized, this will likely be the first and only definitive book on Terracotta by Terracotta team led by Terracotta CTO. The book contains several pragmatic real-world case studies. These empower the reader to build highly scalable, optimized performing enterprise Java applications for financial and even gaming applications. Terracotta is now available in open source options at Terracotta.org.

Mastering Microservices with Java 9 Packt Publishing Ltd

The Comprehensive, Proven Approach to IT Scalability—Updated with New Strategies, Technologies, and Case Studies In *The Art of Scalability*, Second Edition, leading scalability consultants Martin L. Abbott and Michael T. Fisher cover everything you need to know to smoothly scale products and services for any requirement. This extensively revised edition reflects new technologies, strategies, and lessons, as well as new case studies from the authors' pioneering consulting practice, AKF Partners. Writing for technical and nontechnical decision-makers, Abbott and Fisher cover everything that impacts scalability, including architecture, process, people, organization, and technology. Their insights and recommendations reflect more than thirty years of experience at companies ranging from eBay to Visa, and Salesforce.com to Apple. You'll find updated strategies for structuring organizations to maximize agility and scalability, as well as new insights into the cloud (IaaS/PaaS) transition, NoSQL, DevOps, business metrics, and more. Using this guide's tools and advice, you can systematically clear away obstacles to scalability—and achieve unprecedented IT and business performance. Coverage includes • Why scalability problems start with organizations and people, not technology, and what to do about it • Actionable lessons from real successes and failures • Staffing, structuring, and leading the agile, scalable organization • Scaling processes for hyper-growth environments • Architecting scalability: proprietary models for clarifying needs and making choices—including 15 key success principles • Emerging technologies and challenges: data cost, datacenter planning, cloud evolution, and customer-aligned monitoring • Measuring availability, capacity, load, and performance

Practical Node.js Apress

Orchestrate data architecting solutions using Java and related technologies to evaluate, recommend and present the most suitable solution to leadership and clients Key Features Learn how to adapt to the ever-evolving data architecture technology landscape Understand how to choose the best suited technology, platform, and architecture to realize effective business value Implement effective data security and governance principles Book Description Java architectural patterns and tools help architects to build reliable, scalable, and secure data engineering solutions that collect, manipulate, and publish data. This book will help you make the most of the architecting data solutions available with clear and actionable advice from an expert. You'll start with an overview of data architecture, exploring responsibilities of a Java data architect, and learning about various data formats, data storage, databases, and data application platforms as well as how to choose them. Next, you'll understand how to architect a batch and real-time data processing pipeline. You'll also get to grips with the various Java data processing patterns, before progressing to data security and governance. The later chapters will show you how to publish Data as a Service and how you can architect it. Finally, you'll focus on how to evaluate and recommend an architecture by developing performance benchmarks, estimations, and various decision metrics. By the end of this book, you'll be able to successfully orchestrate data architecture solutions using Java and related technologies as well as to evaluate and present the most suitable solution to your clients. What you will learn Analyze and use the best data architecture patterns for problems Understand when and how to choose Java tools for a data architecture Build batch and real-time data engineering solutions using Java Discover how to apply security and governance to a solution Measure performance, publish benchmarks, and optimize solutions Evaluate, choose, and present the best architectural alternatives Understand how to publish Data as a Service using GraphQL and a REST API Who this book is for Data architects, aspiring data architects, Java developers and anyone who wants to develop or optimize scalable data architecture solutions using Java will find this book useful. A basic understanding of data architecture and Java programming is required to get the best from this book.

Pro Java EE 5 Performance Management and Optimization Simon and Schuster

The software architecture landscape has evolved dramatically over the past decade. Microservices have displaced monoliths. Data and applications are increasingly becoming distributed and decentralised. But composing disparate systems is a hard problem. More recently, software practitioners have been rapidly converging on event-driven architecture as a sustainable way of dealing with complexity - integrating systems without increasing their coupling. In *Effective Kafka*, Emil Koutanov explores the fundamentals of Event-Driven Architecture - using Apache Kafka - the world's most popular and supported open-source event streaming platform. You'll learn: - The fundamentals of event-driven architecture and event streaming platforms- The background and rationale behind Apache Kafka, its numerous potential uses and applications- The architecture and core concepts - the underlying software components, partitioning and parallelism, load-balancing, record ordering and consistency modes- Installation of Kafka and related tooling - using standalone deployments, clusters, and containerised deployments with Docker- Using CLI tools to interact with and administer Kafka classes, as well as publishing data and browsing topics- Using third-party web-based tools for monitoring a cluster and gaining insights into the event streams- Building stream processing applications in Java 11 using off-the-shelf client libraries- Patterns and best-practice for organising the application architecture, with emphasis on maintainability and testability of the resulting code- The numerous gotchas that lurk in Kafka's client and broker configuration, and how to counter them- Theoretical background on distributed and concurrent computing, exploring factors affecting their liveness and safety- Best-practices for running multi-tenanted clusters across diverse engineering teams, how teams collaborate to build complex systems at scale and equitably share the cluster with the aid of quotas- Operational aspects of running Kafka clusters at scale, performance tuning and methods for optimising network and storage utilisation- All aspects of Kafka security -including network segregation, encryption, certificates, authentication and

authorization. The coverage is progressively delivered and carefully aimed at giving you a journey-like experience into becoming proficient with Apache Kafka and Event-Driven Architecture. The goal

is to get you designing and building applications. And by the conclusion of this book, you will be a confident practitioner and a Kafka evangelist within your organisation - wielding the knowledge necessary to teach others.

Best Sellers - Books :

- [Fahrenheit 451](#)
- [Rich Dad Poor Dad: What The Rich Teach Their Kids About Money That The Poor And Middle Class Do Not! By Robert T. Kiyosaki](#)
- [Twisted Hate \(twisted, 3\)](#)
- [The Summer I Turned Pretty \(summer I Turned Pretty, The\)](#)
- [The Boy, The Mole, The Fox And The Horse By Charlie Mackesy](#)
- [Young Forever: The Secrets To Living Your Longest, Healthiest Life \(the Dr. Hyman Library, 11\) By Dr. Mark Hyman Md](#)
- [Meditations: A New Translation](#)
- [Stop Overthinking: 23 Techniques To Relieve Stress, Stop Negative Spirals, Declutter Your Mind, And Focus On The Present \(the Path To Calm\) By Nick Trenton](#)
- [Reminders Of Him: A Novel](#)
- [The Collector: A Novel By Daniel Silva](#)