

Best Practices Running Datastax Enterprise Within Docker

Create your own massively scalable Cassandra database with highly responsive database queries About This Book Create a Cassandra cluster and tweak its configuration to get the best performance based on your environment Analyze the key concepts and architecture of Cassandra, which are essential to create highly responsive Cassandra databases A fast-paced and step-by-step guide on handling huge amount of data and getting the best out of your database applications Who This Book Is For If you are a developer who is working with Cassandra and you want to deep dive into the core concepts and understand Cassandra's non-relational nature, then this book is for you. A basic understanding of Cassandra is expected. What You Will Learn Install and set up your Cassandra Cluster using various installation types Use Cassandra Query Language (CQL) to design Cassandra database and tables with various configuration options Design your Cassandra database to be evenly loaded with the lowest read/write latencies Employ the available Cassandra tools to monitor and maintain a Cassandra cluster Debug CQL queries to discover why they are performing relatively slowly Choose the best-suited compaction strategy for your database based on your usage pattern Tune Cassandra based on your deployment operation system environment In Detail Apache Cassandra Essentials takes you step-by-step from from the basics of installation to advanced installation options and database design techniques. It gives you all the information you need to effectively design a well distributed and high performance database. You'll get to know about the steps that are performed by a Cassandra node when you execute a read/write query, which is essential to properly maintain of a Cassandra cluster and to debug any issues. Next, you'll discover how to integrate a Cassandra driver in your applications and perform read/write operations. Finally, you'll learn about the various tools provided by Cassandra for serviceability aspects such as logging, metrics, backup, and recovery. Style and approach This step-by-step guide is packed with examples that explain the core concepts as well as advanced concepts, techniques, and usages of Apache Cassandra.

There are 3 reasons to strengthen your cloud skills: 1- Cloud roles pay well. The average base salary for a Solutions Architect in the U.S. is \$140,000. 2- Cloud skills are in demand. Cloud computing has been one of the most in-demand hard skills for 7 years running. 3- Learning cloud can get you a raise. The average raise received by IT pros who gained new skills and/or certifications is \$15 - 30K. AWS certification is globally recognized as the premier way to demonstrate your AWS cloud skills. The AWS Certified Solutions Architect - Associate Level (SAA-C03) exam validates your ability to effectively demonstrate knowledge of how to architect and deploy secure and robust applications on AWS technologies. It is a required exam for the AWS Certified Solutions Architect - Professional Level certification. In order to prepare for this exam, We suggest purchasing our AWS Certified Solutions Architect - Associate Level Exam Preparation eBook. This AWS Cloud Solutions Architect Associates Certification App covers all of the key concepts you need to know for the AWS Solutions Architect Associate Exam. Solution architecture is a practice of defining and describing an architecture of a system delivered in context of a specific solution and as such it may encompass description of an entire system or only its specific parts. Definition of a solution architecture is typically led by a solution architect. The AWS Certified Solutions Architect - Associate (SAA, SAA-C03) exam is intended for individuals who perform in a solutions architect role. The exam validates a candidate's ability to use AWS technologies to design solutions based on the AWS Well-Architected Framework including: Design solutions that incorporate AWS services to meet current business requirements and future projected needs Design architectures that are secure, resilient, high-performing, and cost-optimized Review existing solutions and determine improvements Become stronger in your current role or prepare to step into a new one by continuing to build the cloud solutions architecture skills companies are begging for right now. Demand for cloud solutions architect proficiency is only set to increase, so you can expect to see enormous ROI on any cloud learning efforts you embark on. What will you learn in this book? Design Secure Architectures Design Resilient Architectures Design High-Performing Architectures Design Cost-Optimized Architectures What are the requirements or prerequisites for reading this book? The target candidate should have at least 1 year of hands-on experience designing cloud solutions that use AWS services Who is this book for? IT Professionals, Solutions Architect, Cloud enthusiasts, Computer Science and Engineering Students, AWS Cloud Developer, Technology Manager and Executives, IT Project Managers What is taught in this book? AWS Certification Preparation for Solutions Architecture - Associate Level Key tools, technologies, and concepts covered • Compute • Cost management • Database • Disaster recovery • High performance • Management and governance • Microservices and component decoupling • Migration and data transfer • Networking, connectivity, and content delivery • Resiliency • Security • Serverless and event-driven design principles • Storage Some New AWS services covered: AWS Data Exchange, AWS Data Pipeline, AWS Lake Formation, Amazon Managed Streaming for Apache Kafka, Amazon AppFlow, AWS Outposts, VMware Cloud on AWS, AWS Wavelength, Amazon Neptune, Amazon Quantum Ledger Database, Amazon Timestream, AWS Amplify, Amazon Comprehend, Amazon Forecast, Amazon Fraud Detector, Amazon Kendra, AWS License Manager, Amazon Managed Grafana, Amazon Managed Service for Prometheus, AWS Proton, Amazon Elastic Transcoder, Amazon Kinesis Video Streams, AWS Application Discovery Service, AWS WAF Serverless, AWS AppSync, etc. Mobile App Version of the AWS Solutions Architect Associates SAA-C03 Certification Exam Prep Book: Android:

https://play.google.com/store/apps/details?id=com.awssolutionarchitectassociatexamprepro.app iOS: https://apps.apple.com/ca/app/solution-architect-assoc-pro/id1501465417 Windows 10/11: https://www.microsoft.com/en-ca/store/p/aws-cert-solution-architect-associate-prep-pro/9pcn58wdr1qr #AWS #SAAC02 #SAAC03 #AWSSAA #SAA #AWSCertification #AWSTraining #LearnAWS #CloudArchitect #SolutionsArchitect #Djmgatech --

The best-selling author of Big Data is back, this time with a unique and in-depth insight into how specific companies use big data. Big data is on the tip of everyone's tongue. Everyone understands its power and importance, but many fail to grasp the actionable steps and resources required to utilise it effectively. This book fills the knowledge gap by showing how major companies are using big data every day, from an up-close, on-the-ground perspective. From technology, media and retail, to sport teams, government agencies and financial institutions, learn the actual strategies and processes being used to learn about customers, improve manufacturing, spur innovation, improve safety and so much more. Organised for easy dip-in navigation, each chapter follows the same structure to give you the information you need quickly. For each company profiled, learn what data was used, what problem it solved and the processes put it place to make it practical, as well as the technical details, challenges and lessons learned from each unique scenario. Learn how predictive analytics helps Amazon, Target, John Deere and Apple understand their customers Discover how big data is behind the success of Walmart, LinkedIn, Microsoft and more Learn how big data is changing medicine, law enforcement, hospitality, fashion, science and banking Develop your own big data strategy by accessing additional reading materials at the end of each chapter

With 28 new chapters, the third edition of The Practice of System and Network Administration innovates yet again! Revised with thousands of updates and clarifications based on reader feedback, this new edition also incorporates DevOps strategies even for non-DevOps environments. Whether you use Linux, Unix, or Windows, this new edition describes the essential practices previously handed down only from mentor to protégé. This wonderfully lucid, often funny cornucopia of information introduces beginners to advanced frameworks valuable for their entire career, yet is structured to help even experts through difficult projects. Other books tell you what commands to type. This book teaches you the cross-platform strategies that are timeless! DevOps techniques: Apply DevOps principles to enterprise IT infrastructure, even in environments without developers Game-changing strategies: New ways to deliver results faster with less stress Fleet management: A comprehensive guide to managing your fleet of desktops, laptops, servers and mobile devices Service management: How to design, launch, upgrade and migrate services Measurable improvement: Assess your operational effectiveness; a forty-page, pain-free assessment system you can start using today to raise the quality of all services Design guides: Best practices for networks, data centers, email, storage, monitoring, backups and more Management skills: Organization design, communication, negotiation, ethics, hiring and firing, and more Have you ever had any of these problems? Have you been surprised to discover your backup tapes are blank? Ever spent a year launching a new service only to be told the users hate it? Do you have more incoming support requests than you can handle? Do you spend more time fixing problems than building the next awesome thing? Have you suffered from a botched migration of thousands of users to a new service? Does your company rely on a computer that, if it died, can't be rebuilt? Is your network a fragile mess that breaks any time you try to improve it? Is there a periodic “hell month” that happens twice a year? Twelve times a year? Do you find out about problems when your users call you to complain? Does your corporate “Change Review Board” terrify you? Does each division of your company have their own broken way of doing things? Do you fear that automation will replace you, or break more than it fixes? Are you underpaid and overworked? No vague “management speak” or empty platitudes. This comprehensive guide provides real solutions that prevent these problems and more!

Monolith to Microservices

Cassandra: The Definitive Guide, (Revised) Third Edition

Applying Graph Thinking and Graph Technologies to Solve Complex Problems

Patterns and Paradigms for Scalable, Reliable Services

A Brief Guide to the Emerging World of Polyglot Persistence

NoSQL Distilled

WINNER OF THE COMEDY WOMEN IN PRINT PRIZE 2020 'Simmonds is a copper-bottomed genius... she is as brilliant a writer as Britain has' Jenny Colgan, Mail Online Cassandra Darke is an art dealer, mean, selfish, solitary by nature, living in Chelsea in a house worth £7 million. She has become a social pariah, but doesn't much care. Between one Christmas and the next, she has sullied the reputation of a West End gallery and has acquired a conviction for fraud, a suspended sentence and a bank balance drained by lawsuits. On the scale of villainy, fraud seems to Cassandra a rather paltry offence - her own crime involving 'no violence, no weapon, no dead body'. But in Cassandra's basement, her young ex-lodger, Nicki, has left a surprise, something which implies at least violence and probably a body . . . Something which forces Cassandra out of her rich enclave and onto the streets. Not those local streets paved with gold and lit with festive glitter, but grimmer, darker places, where she must make the choice between self-sacrifice and running for her life.

Learn Azure in a Month of Lunches, Second Edition, is a tutorial on writing, deploying, and running applications in Azure. In it, you'll work through 21 short lessons that give you real-world experience. Each lesson includes a hands-on lab so you can try out and lock in your new skills. Summary You can be incredibly productive with Azure without mastering every feature, function, and service. Learn Azure in a Month of Lunches, Second Edition gets you up and running quickly, teaching you the most important concepts and tasks in 21 practical bite-sized lessons. As you explore the examples, exercises, and labs, you'll pick up valuable skills immediately and take your first steps to Azure mastery! This fully revised new edition covers core changes to the Azure UI, new Azure features, Azure containers, and the upgraded Azure Kubernetes Service. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Microsoft Azure is vast and powerful, offering virtual servers, application templates, and prebuilt services for everything from data storage to AI. To navigate it all, you need a trustworthy guide. In this book, Microsoft engineer and Azure trainer Iain Foulds focuses on core skills for creating cloud-based applications. About the book Learn Azure in a Month of Lunches, Second Edition, is a tutorial on writing, deploying, and running applications in Azure. In it, you'll work through 21 short lessons that give you real-world experience. Each lesson includes a hands-on lab so you can try out and lock in your new skills. What's inside Understanding Azure beyond point-and-click Securing applications and data Automating your environment Azure services for machine learning, containers, and more About the reader This book is for readers who can write and deploy simple web or client/server applications. About the author Iain Foulds is an engineer and senior content developer with Microsoft. Table of Contents PART 1 - AZURE CORE SERVICES 1 Before you begin 2 Creating a virtual machine 3 Azure Web Apps 4 Introduction to Azure Storage 5 Azure Networking basics PART 2 - HIGH AVAILABILITY AND SCALE 6 Azure Resource Manager 7 High availability and redundancy 8 Load-balancing applications 9 Applications that scale 10 Global databases with Cosmos DB 11 Managing network traffic and routing 12 Monitoring and troubleshooting PART 3 - SECURE BY DEFAULT 13 Backup, recovery, and replication 14 Data encryption 15 Securing information with Azure Key Vault 16 Azure Security Center and updates PART 4 - THE COOL STUFF 17 Machine learning and artificial intelligence 18 Azure Automation 19 Azure containers 20 Azure and the Internet of Things 21 Serverless computing

A #1 New York Times Bestseller! The Shadowhunters must catch a killer in Edwardian London in this dangerous and romantic sequel to the #1 New York Times bestselling novel Chain of Gold, from New York Times and USA TODAY bestselling author Cassandra Clare. Chain of Iron is a Shadowhunters novel. Cordelia Carstairs seems to have everything she ever wanted. She's engaged to marry James Herondale, the boy she has always loved; she has a new life in London with her best friend Lucie; and she bears the sword Cortana, a legendary hero's blade. But the truth is far grimmer. Cordelia's marriage is a lie, arranged to save her reputation, while James remains in love with the Grace Blackthorn. Cortana burns her when she touches it. And a serial murderer is targeting the Shadowhunters of London, killing under cover of darkness, then vanishing without a trace. Now Cordelia, James, and Lucie must follow the trail of the killer through the city's most dangerous streets. All the while, each is keeping a shocking secret: Lucie, that she is attempting to raise the dead; Cordelia, that she has sworn a dangerous oath of loyalty to a mysterious power; and James, that he himself may be the killer they seek.

Black Market Medicine A must read for anyone concerned about the future of healthcare in the U.S. Only a novel so far, John Macklin's story could be yours as a future caregiver. Or you could be cut off from further medical care as your apportionment runs out like Janis. This prescient story carries the reader through a futuristic nightmare where the U.S. Government makes every medical decision for you.

Here is an excerpt: Now she knew there was no God. Janis had long been a believer, but no more. Her body had failed her, and so had her God. For three days, she had asked God to take her from this life. Three long torturous days, lled only with the hell of pain so intense she felt at times she must be losing consciousness. The cancer had spread. There was no hope and not the slightest given from Caregiver No.

42. A glass of water, a wet rag on her forehead, even some word of prayer would have been something. She knew they had cut back her allotment of medicine. She had heard that, as late as 2018, there had still been some Wellness Spas that had provided spiritual assistance. Naturally, Janis knew she shouldn't even be thinking such foolish thoughts, but her mother had laid the seed of spiritualism in her many years before. Before the legislation of 2045. Black Market Medicine Mainstream Fiction By Cassandra Collins facebook.com/BlackMarketMedicine

An Authoritative Guide to Building Microservices, Web and Enterprise Applications, and Best Practices

Cassandra Design Patterns

Strategies and Best Practices for Connected Products and Services

250+ AWS SAA SAA-C02 SAA-C03 Quizzes, Practice Exams, Cheat Sheets, Flashcards, I passed SAA Testimonials, Tips

Briggs

Hearing Before the Committee on Small Business, United States Senate, One Hundred Fourth Congress, First Session, April 4, 1995

Efficient Processing of Deep Neural Networks

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: 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*

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

Data is bigger, arrives faster, and comes in a variety of formats—and it all needs to be processed at scale for analytics or machine learning. But how can you process such varied workloads efficiently? Enter Apache Spark. Updated to include Spark 3.0, this second edition shows data engineers and data scientists why structure and unification in Spark matters. Specifically, this book explains how to perform simple and complex data analytics and employ machine learning algorithms. Through step-by-step walk-throughs, code snippets, and notebooks, you'll be able to: Learn Python, SQL, Scala, or Java high-level Structured APIs Understand Spark operations and SQL Engine Inspect, tune, and debug Spark operations with Spark configurations and Spark UI Connect to data sources: JSON, Parquet, CSV, Avro, ORC, Hive, S3, or Kafka Perform analytics on batch and streaming data using Structured Streaming Build reliable data pipelines with open source Delta Lake and Spark Develop machine learning pipelines with MLlib and productionize models using MLflow

Data Science & Business Analytics explores the application of big data and business analytics by academics, researchers, industrial experts, policy makers and practitioners, helping the reader to understand how big data can be efficiently utilized in better managerial applications.

Distributed Data at Web Scale

A Developer's Approach

Cloud Native Transformation

Trino: The Definitive Guide

Design Patterns for Cloud Native Applications

Cassandra Darke

Streamlined Enterprise Data Management and Analysis

"An Airflow bible. Useful for all kinds of users, from novice to expert." - Rambabu Posa, Sai Aashika Consultancy Data Pipelines with Apache Airflow teaches you how to build and maintain effective data pipelines. A successful pipeline moves data efficiently, minimizing pauses and blockages between tasks, keeping every process along the way operational. Apache Airflow managing data pipelines, eliminating the need for a hodgepodge collection of tools, snowflake code, and homegrown processes. Using real-world scenarios and examples, Data Pipelines with Apache Airflow teaches you how to simplify and automate data pipelines, reduce operational overhead, and smoothly integrate all the technologies in your stack. Purchase of the book from Manning Publications. About the technology Data pipelines manage the flow of data from initial collection through consolidation, cleaning, analysis, visualization, and more. Apache Airflow provides a single platform you can use to design, implement, monitor, and maintain your pipelines. Its easy-to-use UI, plug-and-play options, and flexible Python scripting make it a breeze to use. Data Pipelines with Apache Airflow teaches you how to build and maintain effective data pipelines. You'll explore the most common usage patterns, including aggregating multiple data sources, connecting to and from data lakes, and cloud deployment. Part reference and part tutorial, this practical guide covers every aspect of the directed acyclic graphs (DAGs) that power Airflow. What's inside Build, test, and deploy Airflow pipelines as DAGs Automate moving and transforming data Analyze historical datasets using backfilling Develop custom components Set up Airflow in production environments About the reader For DevOps, data engineers, machine learning engineers, and sysadmins with intermediate Python skills. About the author Bas Harshani has extensive experience using Airflow to develop pipelines for major companies. Bas is also an Airflow committer. Table of Contents PART 1 - GETTING STARTED 1 Meet Apache Airflow 2 Anatomy of an Airflow DAG 3 Scheduling in Airflow 4 Templating tasks using the Airflow context 5 Defining dependencies between tasks PART 2 - BEYOND THE BASICS 6 Triggering workflows 7 Cloud deployment 8 Components 9 Testing 10 Running tasks in containers PART 3 - AIRFLOW IN PRACTICE 11 Best practices 12 Operating Airflow in production 13 Securing Airflow 14 Project: Finding the fastest way to get around NYC PART 4 - IN THE CLOUDS 15 Airflow in the clouds 16 Airflow on AWS 17 Airflow on Azure 18 Airflow in GCP

Graph data closes the gap between the way humans and computers view the world. While computers rely on static rows and columns of data, people navigate and reason about life through relationships. This practical guide demonstrates how graph data brings these two approaches together. By working with concepts from graph theory, database schema, distributed systems, and known as graph thinking. Authors Denise Koessler Gosnell and Matthias Broecheler show data engineers, data scientists, and data analysts how to solve complex problems with graph databases. You'll explore templates for building with graph technology, along with examples that demonstrate how teams think about graph data within an application. Build an example application using graph technology to build a Customer 360 application, the most popular graph data pattern today Dive into hierarchical data and troubleshoot a new paradigm that comes from working with graph data Find paths in graph data and learn why your trust in different paths motivates and informs your preferences Use collaborative filtering to design a Netflix-inspired recommendation system This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high energy consumption. enable efficient processing of deep neural networks to improve metrics—such as energy-efficiency, throughput, and latency—without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing: a description and taxonomy of hardware architectural approaches, comparing different designs; features of the DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as a formalization and organization of key concepts from contemporary works that provides insight into the challenges. 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 revised third edition—updated for Cassandra 4.0 and new developments in the Cassandra ecosystem, including deployment and practical examples to help you put this database to work in a production environment. Authors Jeff Carpenter and Eben Hewitt demonstrate the advantages of Cassandra's nonrelational design, with special attention to data modeling. Developers, DBAs, and application architects looking to solve a database scaling issue or future-proof an application will learn how to design and deploy Cassandra's distributed and decentralized structure Use the Cassandra Query Language (CQL) and cqsh (the CQL shell) Create a working data model and compare it with an equivalent relational model Design and develop applications using client drivers Explore cluster topology and learn how nodes exchange data Maintain a high level of performance in your cluster

Kubernetes Integrate Cassandra with Spark, Kafka, Elasticsearch, Solr, and Lucene
Microservice Architecture
Enterprise Data Workflows with Cascading

Apache Cassandra Essentials
Practical Cassandra

Big Data in Practice
The Small Business Administration's 8(a) Minority Business Development Program

AWS Solutions Architect Associates SAA-CO3 Certification Practice Tests and Quizzes illustrated

The need to handle increasingly larger data volumes is one factor driving the adoption of a new class of nonrelational “NoSQL” databases. Advocates of NoSQL databases claim they can be used to build systems that are more performant, scale better, and are easier to program. NoSQL Distilled is a concise but thorough introduction to this rapidly emerging technology. Pramod J. Sadalage and Martin Fowler explain how NoSQL databases work and the ways that they may be a superior alternative to a traditional RDBMS. The authors provide a fast-paced guide to the concepts you need to know in order to evaluate whether NoSQL databases are right for your needs and, if so, which technologies you should explore further. The first part of the book concentrates on core concepts, including schemaless data models, aggregates, new distribution models, the CAP theorem, and map-reduce. In the second part, the authors explore architectural and design issues associated with implementing NoSQL. They also present realistic use cases that demonstrate NoSQL databases at work and feature representative examples using Riak, MongoDB, Cassandra, and Neo4j. In addition, by drawing on Pramod Sadalage's pioneering work, NoSQL Distilled shows how to implement evolutionary design with schema migration: an essential technique for applying NoSQL databases. The book concludes by describing how NoSQL is ushering in a new age of Polyglot Persistence, where multiple data-storage worlds coexist, and architects can choose the technology best optimized for each type of data access.

How do you start? How should you build a plan for cloud migration for your entire portfolio? How will your organization be affected by these changes? This book, based on real-world cloud experiences by enterprise IT teams, seeks to provide the answers to these questions. Here, you'll see what makes the cloud so compelling to enterprises; with which applications you should start your cloud journey; how your organization will change, and how skill sets will evolve; how to measure progress; how to think about security, compliance, and business buy-in; and how to exploit the ever-growing feature set that the cloud offers to gain strategic and competitive advantage.

If you create, manage, operate, or configure systems running in the cloud, you're a cloud engineer—even if you work as a system administrator, software developer, data scientist, or site reliability engineer. With this book, professionals from around the world provide valuable insight into today's cloud engineering role. These concise articles explore the entire cloud computing experience, including fundamentals, architecture, and migration. You'll delve into security and compliance, operations and reliability, and software development. And examine networking, organizational culture, and more. You're sure to find 1, 2, or 97 things that inspire you to dig deeper and expand your own career. "Three Keys to Making the Right Multicloud Decisions," Brendan O'Leary "Serverless Bad Practices," Manases Jesus Galindo Bello "Failing a Cloud Migration," Lee Atchison "Treat Your Cloud Environment as If It Were On Premises," Iyana Garry "What Is Toil, and Why Are SREs Obsessed with It?," Zachary Nickens "Lean QA: The QA Evolving in the DevOps World," Theresa Neate "How Economies of Scale Work in the Cloud," Jon Moore "The Cloud Is Not About the Cloud," Ken Corless "Data Gravity: The Importance of Data Management in the Cloud," Geoff Hughes "Even in the Cloud, the Network Is the Foundation," David Murray "Cloud Engineering Is About Culture, Not Containers," Holly Cummins

"Eric and Russell were early adopters of Cassandra at SimpleReach. In Practical Cassandra, you benefit from their experience in the trenches administering Cassandra, developing against it, and building one of the first CQL drivers. If you are deploying Cassandra soon, or you inherited a Cassandra cluster to tend, spend some time with the deployment, performance tuning, and maintenance chapters.. If you are new to Cassandra, I highly recommend the chapters on data modeling and CQL." –From the Foreword by Jonathon Ellis, Apache Cassandra Chair Build and Deploy Massively Scalable, Super-fast Data Management Applications with Apache Cassandra Practical Cassandra is the first hands-on developer's guide to building Cassandra systems and applications that deliver breakthrough speed, scalability, reliability, and performance. Fully up to date, it reflects the latest versions of Cassandra—including Cassandra Query Language (CQL), which dramatically lowers the learning curve for Cassandra developers. Pioneering Cassandra developers and Datastax MVPs Russell Bradberry and Eric Lubow walk you through every step of building a real production application that can store enormous amounts of structured, semi-structured, and unstructured data. Drawing on their exceptional expertise, Bradberry and Lubow share practical insights into issues ranging from querying to deployment, management, maintenance, monitoring, and troubleshooting. The authors cover key issues, from architecture to migration, and guide you through crucial decisions about configuration and data modeling. They provide tested sample code, detailed explanations of how Cassandra works “under the covers,” and new case studies from three cutting-edge users: Ooyala, Hailo, and eBay. Coverage includes Understanding Cassandra's approach, architecture, key concepts, and primary use cases— and why it's so blazingly fast Getting Cassandra up and running on single nodes and large clusters Applying the new design patterns, philosophies, and features that make Cassandra such a powerful data store Leveraging CQL to simplify your transition from SQL-based RDBMSes Deploying and provisioning through the cloud or on bare-metal hardware Choosing the right configuration options for each type of workload Tweaking Cassandra to get maximum performance from your hardware, OS, and JVM Mastering Cassandra's essential tools for maintenance and monitoring Efficiently solving the most common problems with Cassandra deployment, operation, and application development

Data Pipelines with Apache Airflow

Design, develop, and deploy highly adaptable, scalable, and secure RESTful web APIs

Data Science & Business Analytics

Enterprise Cloud ePUB 1

Hands-On RESTful API Design Patterns and Best Practices

Serverless Design Patterns and Best Practices

Expert Apache Cassandra Administration

Quickly and productively develop complex Spring applications and microservices out of the box, with minimal concern over things like configurations. This revised book will show you how to fully leverage the Spring Boot 2 technology and how to apply it to create enterprise ready applications that just work. It will also cover what's been added to the new Spring Boot 2 release, including Spring Framework 5 features like WebFlux, Security, Actuator and the new way to expose Metrics through Micrometer framework, and more. This book is your authoritative hands-on practical guide for increasing your enterprise Java and cloud application productivity while decreasing development time. It's a no nonsense guide with case studies of increasing complexity throughout the book. The author, a senior solutions architect and Principal Technical instructor with Pivotal, the company behind the Spring Framework, shares his experience, insights and first-hand knowledge about how Spring Boot technology works and best practices. Pro Spring Boot 2 is an essential book for your Spring learning and reference library. What You Will Learn Configure and use Spring Boot Use non-functional requirements with Spring Boot Actuator Carry out web development with Spring Boot Persistence with JDBC, JPA and NoSQL Databases Messaging with JMS, RabbitMQ and WebSockets Test and deploy with Spring Boot A quick look at the Spring Cloud projects Microservices and deployment to the Cloud Extend Spring Boot by creating your own Spring Boot Starter and @Enable feature Who This Book Is For Experienced Spring and Java developers seeking increased productivity gains and decreased complexity and development time in their applications and software services.

Legend has it that Google deploys over two billion application containers a week. How's that possible? Google revealed the secret through a project called Kubernetes, an open source cluster orchestrator (based on its internal Borg system) that radically simplifies the task of building, deploying, and maintaining scalable distributed systems in the cloud. This practical guide shows you how Kubernetes and container technology can help you achieve new levels of velocity, agility, reliability, and efficiency. Authors Kelsey Hightower, Brendan Burns, and Joe Beda—who've worked on Kubernetes at Google and other organizations—explain how this system fits into the lifecycle of a distributed application. You will learn how to use tools and APIs to automate scalable distributed systems, whether it is for online services, machine-learning applications, or a cluster of Raspberry Pi computers. Explore the distributed system challenges that Kubernetes addresses Dive into containerized application development, using containers such as Docker Create and run containers on Kubernetes, using the docker image format and container runtime Explore specialized objects essential for running applications in production Reliably roll out new software versions without downtime or errors Get examples of how to develop and deploy real-world applications in Kubernetes

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.

Get started with designing your serverless application using optimum design patterns and industry standard practices Key Features Learn the details of popular software patterns and how they are applied to serverless applications Understand key concepts and components in serverless designs Walk away with a thorough understanding of architecting serverless applications Book Description Serverless applications handle many problems that developers face when running systems and servers. The serverless pay-per-invocation model can also result in drastic cost savings, contributing to its popularity. While it's simple to create a basic serverless application, it's critical to structure your software correctly to ensure it continues to succeed as it grows. Serverless Design Patterns and Best Practices presents patterns that can be adapted to run in a serverless environment. You will learn how to develop applications that are scalable, fault tolerant, and well-tested. The book begins with an introduction to the different design pattern categories available for serverless applications. You will learn the trade-offs between GraphQL and REST and how they fare regarding overall application design in a serverless ecosystem. The book will also show you how to migrate an existing API to a serverless backend using AWS API Gateway. You will learn how to build event-driven applications using queuing and streaming systems, such as AWS Simple Queuing Service (SQS) and AWS Kinesis. Patterns for data-intensive serverless application are also explained, including the lambda architecture and MapReduce. This book will equip you with the knowledge and skills you need to develop scalable and resilient serverless applications confidently. What you will learn Comprehend the popular design patterns currently being used with serverless architectures Understand the various design options and corresponding implementations for serverless web application APIs Learn multiple patterns for data-intensive serverless systems and pipelines, including MapReduce and Lambda Architecture Learn how to leverage hosted databases, queues, streams, storage services, and notification services Understand error handling and system monitoring in a serverless architecture a serverless architecture Learn how to set up a serverless application for continuous integration, continuous delivery, and continuous deployment Who this book is for If you're a software architect, engineer, or someone who wants to build serverless applications, which are non-trivial in complexity and scope, then this book is for you. Basic knowledge of programming and serverless computing concepts are assumed.

Enterprise IoT

Dive into the Future of Infrastructure

Simca's Cuisine

Get up and running with the fundamentals and functionalities of seven of the most popular NoSQL databases

Build, secure, and deploy enterprise ready serverless applications with AWS to improve developer productivity

The French Chef Cookbook

97 Things Every Cloud Engineer Should Know

How do you detangle a monolithic system and migrate it to a microservice architecture? How do you do it while maintaining business-as-usual? As a companion to Sam Newman's extremely popular Building Microservices, this new book details a proven method for transitioning an existing monolithic system to a microservice architecture. With many illustrative examples, insightful migration patterns, and a bevy of practical advice to transition your monolith enterprise into a microservice operation, this practical guide covers multiple scenarios and strategies for a successful migration, from initial planning all the way through application and database decomposition. You'll learn several tried and tested patterns and techniques that you can use as you migrate your existing architecture. Ideal for organizations looking to transition to microservices, rather than rebuild Helps companies determine whether to migrate, when to migrate, and where to begin Addresses communication, integration, and the migration of legacy systems Discusses multiple migration patterns and where they apply Provides database migration examples, along with synchronization strategies Explores application decomposition, including several architectural refactoring patterns Delves into details of database decomposition, including the impact of breaking referential and transactional integrity, new failure modes, and more

So many classics of French cuisine, such as Boeuf Bourguignon, Quiche Lorraine, and Coq au Vin, owe their presence in America's kitchens to one extraordinary chef: Julia Child. She awakened the American palate to this cuisine, and was devoted to making it accessible to home cooks. In 1963, her personality came to life with her public TV show, ¿The French Chef.¿ The show was a tremendous success and made Julia a household name and a culinary celebrity. This book presents all the classic French recipes premiered on ¿The French Chef.¿ Julia takes the reader through the essential

techniques of this cuisine. Her notes on ingred. and equipment and her step-by-step directions place every dish within reach of home cooks. Drawings and photos.

Build real-world, industry-strength data storage solutions with time-tested design methodologies using Cassandra About This Book Explore design patterns which co-exist with legacy data stores, migration from RDBMS, and caching technologies with Cassandra Learn about design patterns and use Cassandra to provide consistency, availability, and partition tolerance guarantees for applications Handle temporal data for analytical purposes Who This Book Is For This book is intended for big data developers who are familiar with the basics of Cassandra and wish to understand and utilize Cassandra design patterns to develop real-world big data solutions. Prior knowledge of RDBMS solutions is assumed. What You Will Learn Enable Cassandra to co-exist with RDBMS and other legacy data stores Explore various design patterns to build effective and robust storage solutions Migrate from RDBMS-based data stores and caching solutions to Cassandra Understand the behaviour of Cassandra when trying to balance the needs of consistency, availability, and partition tolerance Deal with time stamps related to data effectively See how Cassandra can be used in analytical use cases Apply the design patterns covered in this book in real-world use cases In Detail There are many NoSQL data stores used by big data applications. Cassandra is one of the most widely used NoSQL data stores that is frequently used by a huge number of heavy duty Internet-scale applications. Unlike the RDBMS world, the NoSQL landscape is very diverse and there is no one way to model data stores. This mandates the need to have good solutions to commonly seen data store design problems. Cassandra addresses such common problems simply. If you are new to Cassandra but well-versed in RDBMS modeling and design, then it is natural to model data in the same way in Cassandra, resulting in poorly performing applications and losing the real purpose of Cassandra. If you want to learn to make the most of Cassandra, this book is for you. This book starts with strategies to integrate Cassandra with other legacy data stores and progresses to the ways in which a migration from RDBMS to Cassandra can be accomplished. The journey continues with ideas to migrate data from cache solutions to Cassandra. With this, the stage is set and the book moves on to some of the most commonly seen problems in applications when dealing with consistency, availability, and partition tolerance guarantees. Cassandra is exceptionally good at dealing with temporal data and patterns such as the time-series pattern and log pattern, which are covered next. Many NoSQL data stores fail miserably when a huge amount of data is read for analytical purposes, but Cassandra is different in this regard. Keeping analytical needs in mind, you'll walk through different and interesting design patterns. No theoretical discussions are complete without a good set of use cases to which the knowledge gained can be applied, so the book concludes with a set of use cases you can apply the patterns you've learned. Style and approach This book is written in very simple language and an engaging style complete with examples in every chapter and real-world use cases at the end of the book.

There is an easier way to build Hadoop applications. With this hands-on book, you'll learn how to use Cascading, the open source abstraction framework for Hadoop that lets you easily create and manage powerful enterprise-grade data processing applications—without having to learn the intricacies of MapReduce. Working with sample apps based on Java and other JVM languages, you'll quickly learn Cascading's streamlined approach to data processing, data filtering, and workflow optimization. This book demonstrates how this framework can help your business extract meaningful information from large amounts of distributed data. Start working on Cascading example projects right away Model and analyze unstructured data in any format, from any source Build and test applications with familiar constructs and reusable components Work with the Scalding and Cascalog Domain-Specific Languages Easily deploy applications to Hadoop, regardless of cluster location or data size Build workflows that integrate several big data frameworks and processes Explore common use cases for Cascading, including features and tools that support them Examine a case study that uses a dataset from the Open Data Initiative

Mastering Kubernetes

Designing Distributed Systems

Learn Azure in a Month of Lunches, Second Edition

Pro Spring Boot 2

Chain of Iron

Aligning Principles, Practices, and Culture

Black Market Medicine

"Eric and Russell were early adopters of Cassandra at SimpleReach. In Practical Cassandra, you benefit from their experience in the trenches administering Cassandra, developing against it, and building one of the first CQL drivers. If you are deploying Cassandra soon, or you inherited a Cassandra cluster to tend, spend some time with the deployment, performance tuning, and maintenance chapters... If you are new to Cassandra, I highly recommend the chapters on data modeling and CQL." –From the Foreword by Jonathon Ellis, Apache Cassandra Chair Build and Deploy Massively Scalable, Super-fast Data Management Applications with Apache Cassandra Practical Cassandra is the first hands-on developer's guide to building Cassandra systems and applications that deliver breakthrough speed, scalability, reliability, and performance. Fully up to date, it reflects the latest versions of Cassandra—including Cassandra Query Language (CQL), which dramatically lowers the learning curve for Cassandra developers.

Pioneering Cassandra developers and Datastax MVPs Russell Bradberry and Eric Lubow walk you through every step of building a real production application that can store enormous amounts of structured, semi-structured, and unstructured data. Drawing on their exceptional expertise, Bradberry and Lubow share practical insights into issues ranging from querying to deployment, management, maintenance, monitoring, and troubleshooting. The authors cover key issues, from architecture to migration, and guide you through crucial decisions about configuration and data modeling. They provide tested sample code, detailed explanations of how Cassandra works "under the covers," and new case studies from three cutting-edge users: Ooyala, Hailo, and eBay. Coverage includes Understanding Cassandra's approach, architecture, key concepts, and primary use cases– and why it's so blazingly fast Getting Cassandra up and running on single nodes and large clusters Applying the new design patterns, philosophies, and features that make Cassandra such a powerful data store Leveraging CQL to simplify your transition from SQL-based RDBMSes Deploying and provisioning through the cloud or on bare-metal hardware Choosing the right configuration options for each type of workload Tweaking Cassandra to get maximum performance from your hardware, OS, and JVM Mastering Cassandra's essential tools for maintenance and monitoring Efficiently solving the most common problems with Cassandra deployment, operation, and application development

Simone ("Simca") Beck is known to millions of Americans as Julia Child's French partner in the creation of the two classic volumes of Mastering the Art of French Cooking. Now, she gives us her own delectable recipes—the ones she most treasures out of a lifetime of cooking creativity that has made her one of the great cuisinières of her day. Here are recipes that were inspired by old French family specialties found in her mother's and grandmother's well-thumbed notebooks: recipes that grew out of Simca's life in the provinces (particularly Normandy, Alsace, and Provence) where she has gardened, cooked, dined out, and entertained: simple delights and fabulous concoctions all set down with a beautiful French clarity. Skillfully adapting her French ways to American needs, she presents over 100 recipes in 31 alluring menus designed for every sort of occasion—a warming dinner after a winter walk in the woods, a feast to dazzle your friends, a buffet for winter and one for summer, even an elegant picnic. For each menu Simca has written a charming, altogether personal introduction filled with nuggets of useful information, like what can be cooked ahead of time or how long last-minute preparations will take. Specific wines are always suggested with the menus, along with specific cheeses when called for. In addition, this volume features a small collection of other favorite dishes that did not fit into the menus but were simply too good to leave out. All in all, Simca's Cuisine is a lasting treasure for everyone in search of new delicacies to serve, new menus that will enchant, new aromas and flavors in the French tradition, and new ways to find expertise in the kitchen and joy at the table.

Follow this handbook to build, configure, tune, and secure Apache Cassandra databases. Start with the installation of Cassandra and move on to the creation of a single instance, and then a cluster of Cassandra databases. Cassandra is increasingly a key player in many big data environments, and this book shows you how to use Cassandra with Apache Spark, a popular big data processing framework. Also covered are day-to-day topics of importance such as the backup and recovery of Cassandra databases, using the right compression and compaction strategies, and loading and unloading data. Expert Apache Cassandra Administration provides numerous step-by-step examples starting with the basics of a Cassandra database, and going all the way through backup and recovery, performance optimization, and monitoring and securing the data. The book serves as an authoritative and comprehensive guide to the building and management of simple to complex Cassandra databases. The book: Takes you through building a Cassandra database from installation of the software and creation of a single database, through to complex clusters and data centers Provides numerous examples of actual commands in a real-life Cassandra environment that show how to confidently configure, manage, troubleshoot, and tune Cassandra databases Shows how to use the Cassandra configuration properties to build a highly stable, available, and secure Cassandra database that always operates at peak efficiency What You'll Learn Install the Cassandra software and create your first database Understand the Cassandra data model, and the internal architecture of a Cassandra database Create your own Cassandra cluster, step-by-step Run a Cassandra cluster on Docker Work with Apache Spark by connecting to a Cassandra database Deploy Cassandra clusters in your data center, or on Amazon EC2 instances Back up and restore mission-critical Cassandra databases Monitor, troubleshoot, and tune production Cassandra databases, and cut your spending on resources such as memory, servers, and storage Who This Book Is For Database administrators, developers, and architects who are looking for an authoritative and comprehensive single volume for all their Cassandra administration needs. Also for administrators who are tasked with setting up and maintaining highly reliable and high-performing Cassandra databases. An excellent choice for big data administrators, database administrators, architects, and developers who use Cassandra as their key data store, to support high volume online transactions, or as a decentralized, elastic data store.

Perform fast interactive analytics against different data sources using the Trino high-performance distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Trino. Initially developed by Facebook, open source Trino is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber, and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Trino query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Trino's use cases and learn about tools that will help you connect to Trino and query data Go deeper: Learn Trino's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Trino in production:

Secure Trino, monitor workloads, tune queries, and connect more applications; learn how other organizations apply Trino

Practical Patterns for Innovation

Learning Spark

The Practitioner's Guide to Graph Data

The Practice of System and Network Administration

How 45 Successful Companies Used Big Data Analytics to Deliver Extraordinary Results

Volume 1: DevOps and other Best Practices for Enterprise IT

Principles and Best Practices of Scalable Realtime Data Systems

Current hype aside, the Internet of Things will ultimately become as fundamental as the Internet itself, with lots of opportunities and trials along the way. To help you navigate these choppy waters, this practical guide introduces a dedicated methodology for businesses preparing to transition towards IoT-based business models. With a set of best practices based on case study analysis, expert interviews, and the authors' own experience, the Ignite | IoT Methodology outlined in this book delivers actionable guidelines to assist you with IoT strategy management and project execution. You'll also find a detailed case study of a project fully developed with this methodology. This book consists of three parts: Illustrative case studies of selected IoT domains, including smart energy, connected vehicles, manufacturing and supply chain management, and smart cities The Ignite | IoT Methodology for defining IoT strategy, preparing your organization for IoT adoption, and planning and executing IoT projects A detailed case study of the IIC Track & Trace testbed, one of the first projects to be fully developed according to the Ignite | IoT Methodology

In the past few years, going cloud native has been a big advantage for many companies. But it's a tough technique to get right, especially for enterprises with critical legacy systems. This practical hands-on guide examines effective architecture, design, and cultural patterns to help you transform your organization into a cloud native enterprise—whether you're moving from older architectures or creating new systems from scratch. By following Wealth Grid, a fictional company, you'll understand the challenges, dilemmas, and considerations that accompany a move to the cloud. Technical managers and architects will learn best practices for taking on a successful company-wide transformation. Cloud migration consultants Pini Reznik, Jamie Dobson, and Michelle Genow draw patterns from the growing community of expert practitioners and enterprises that have successfully built cloud native systems. You'll learn what works and what doesn't when adopting cloud native—including how this transition affects not just your technology but also your organizational structure and processes. You'll learn: What cloud native means and why enterprises are so interested in it Common barriers and pitfalls that have affected other companies (and how to avoid them) Context-specific patterns for a successful cloud native transformation How to implement a safe, evolutionary cloud native approach How companies addressed root causes and misunderstandings that hindered their progress Case studies from real-world companies that have succeeded with cloud native transformations

Cassandra: The Definitive Guide, (Revised) Third Edition"O'Reilly Media, Inc."

In the race to compete in today's fast-moving markets, large enterprises are busy adopting new technologies for creating new products, processes, and business models. But one obstacle on the road to digital transformation is placing too much emphasis on technology, and not enough on the types of processes technology enables. What if different lines of business could build their own services and applications—and decision-making was distributed rather than centralized? This report explores the concept of a digital business platform as a way of empowering individual business sectors to act on data in real time. Much innovation in a digital enterprise will increasingly happen at the edge, whether it involves business users (from marketers to data scientists) or IoT devices. To facilitate the process, your core IT team can provide these sectors with the digital tools they need to innovate quickly. This report explores: Key cultural and organizational changes for developing business capabilities through cross-functional product teams A platform for integrating applications, data sources, business partners, clients, mobile apps, social networks, and IoT devices Creating internal API programs for building innovative edge services in low-code or no-code environments Tools including Integration Platform as a Service, Application Platform as a Service, and Integration Software as a Service The challenge of integrating microservices and serverless architectures Event-driven architectures for processing and reacting to events in real time You'll also learn about a complete pervasive integration solution as a core component of a digital business platform to serve every audience in your organization.

Kubernetes: Up and Running

Cassandra: The Definitive Guide

Big Data

Seven NoSQL Databases in a Week

Evolutionary Patterns to Transform Your Monolith

REST architecture (style) is a pivot of distributed systems, simplify data integration amongst modern and legacy applications leverages through the RESTful paradigm. This book is fully loaded with many RESTful API patterns, samples, hands-on implementations and also discuss the capabilities of many REST API frameworks for Java, Scala, Python and Go

Master the art of container management utilizing the power of Kubernetes. About This Book This practical guide demystifies Kubernetes and ensures that your clusters are always available, scalable, and up to date Discover new features such as autoscaling, rolling updates, resource quotas, and cluster size Master the skills of designing and deploying large clusters on various cloud platforms Who This Book Is For The book is for system administrators and developers who have intermediate level of knowledge with Kubernetes and are now waiting to master its advanced features. You should also have basic networking knowledge. This advanced-level book provides a pathway to master Kubernetes. What You Will Learn Architect a robust Kubernetes cluster for long-time operation Discover the advantages of running Kubernetes on GCE, AWS, Azure, and bare metal See the identity model of Kubernetes and options for cluster federation Monitor and troubleshoot Kubernetes clusters and run a highly available Kubernetes Create and configure custom Kubernetes resources and use third-party resources in your automation workflows Discover the art of running complex stateful applications in your container environment Deliver applications as standard packages In Detail Kubernetes is an open source system to automate the deployment, scaling, and management of containerized applications. If you are running more than just a few containers or want automated management of your containers, you need Kubernetes. This book mainly focuses on the advanced management of Kubernetes clusters. It covers problems that arise when you start using container orchestration in production. We start by giving you an overview of the guiding principles in Kubernetes design and show you the best practises in the fields of security, high availability, and cluster federation. You will discover how to run complex stateful microservices on Kubernetes including advanced features as horizontal pod autoscaling, rolling updates, resource quotas, and persistent storage back ends. Using real-world use cases, we explain the options for network configuration and provides guidelines on how to set up, operate, and troubleshoot various Kubernetes networking plugins. Finally, we cover custom resource development and utilization in automation and maintenance workflows. By the end of this book, you'll know everything you need to know to go from intermediate to advanced level. Style and approach Delving into the design of the Kubernetes platform, the reader will be exposed to the advanced features and best practices of Kubernetes. This book will be an advanced level book which will provide a pathway to master Kubernetes

With the immense cost savings and scalability the cloud provides, the rationale for building cloud native applications is no longer in question. The real issue is how. With this practical guide, developers will learn about the most commonly used design patterns for building cloud native applications using APIs, data, events, and streams in both greenfield and brownfield development. You'll learn how to incrementally design, develop, and deploy large and effective cloud native applications that you can manage and maintain at scale with minimal cost, time, and effort. Authors Kasun Indrasiri and Sriskandarajah Subhothayan highlight use cases that effectively demonstrate the challenges you might encounter at each step. Learn the fundamentals of cloud native applications Explore key cloud native communication, connectivity, and composition patterns Learn decentralized data management techniques Use event-driven architecture to build distributed and scalable cloud native applications Explore the most commonly used patterns for API management and consumption Examine some of the tools and technologies you'll need for building cloud native systems

Microservices can have a positive impact on your enterprise—just ask Amazon and Netflix—but you can fall into many traps if you don't approach them in the right way. This practical guide covers the entire microservices landscape, including the principles, technologies, and methodologies of this unique, modular style of system building. You'll learn about the experiences of organizations around the globe that have successfully adopted microservices. In three parts, this book explains how these services work and what it means to build an application the Microservices Way. You'll explore a design-based approach to microservice architecture with guidance for implementing various elements. And you'll get a set of recipes and practices for meeting practical, organizational, and cultural challenges to microservice adoption. Learn how microservices can help you drive business objectives Examine the principles, practices, and culture that define microservice architectures Explore a model for creating complex systems and a design process for building a microservice architecture Learn the fundamental design concepts for individual microservices Delve into the operational elements of a microservices architecture, including containers and service discovery Discover how to handle the challenges of introducing microservice architecture in your organization