

Red Hat Amqp User Guide

In simple terms, the book is designed to give IT professionals an extensive idea of what cloud computing is all about, the basic fundamentals, what the different options of cloud computing are for an enterprise, and how the same can be adopted to their own enterprise. This book is exhaustive and covers almost all the top cloud computing technologies and to the lowest level of details, which will help even a junior-level IT professional to design and deploy cloud solutions based on the individual requirements. This book offers high level of details, which will help IT administrators to manage and maintain the corporate and SME IT infrastructure. This book can also be a part of an engineering curriculum, especially where information technology and computer science courses are offered.

Summary Camel in Action, Second Edition is the most complete Camel book on the market. Written by core developers of Camel and the authors of the highly acclaimed first edition, this book distills their experience and practical insights so that you can tackle integration tasks like a pro. Forewords by James Strachan and Dr. Mark Little Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Apache Camel is a Java framework that implements enterprise integration patterns (EIPs) and comes with over 200 adapters to third-party systems. A concise DSL lets you build integration logic into your app with just a few lines of Java or XML. By using Camel, you benefit from the testing and experience of a large and vibrant open source community. About the Book Camel in Action, Second Edition is the definitive guide to the Camel framework. It starts with core concepts like sending, receiving, routing, and transforming data. It then goes in depth on many topics such as how to develop, debug, test, deal with errors, secure, scale, cluster, deploy, and monitor your Camel applications. The book also discusses how to run Camel with microservices, reactive systems, containers, and in the cloud. What's Inside Coverage of all relevant EIPs Camel microservices with Spring Boot Camel on Docker and Kubernetes Error handling, testing, security, clustering, monitoring, and deployment Hundreds of examples in Java and XML About the Reader Readers should be familiar with Java. This book is accessible to beginners and invaluable to experts. About the Author Claus Ibsen is a senior principal engineer working for Red Hat specializing in cloud and integration. He has worked on Apache Camel for the last nine years where he heads the project. Claus lives in Denmark. Jonathan Anstey is an engineering manager at Red Hat and a core Camel contributor. He lives in Newfoundland, Canada. Table of Contents Part 1 - First steps Meeting Camel Routing with Camel Part 2 - Core Camel Transforming data with Camel Using beans with Camel Enterprise integration patterns Using components Part 3 - Developing and testing Microservices Developing Camel projects Testing RESTful web services Part 4 - Going further with Camel Error handling Transactions and idempotency Parallel processing Securing Camel Part 5 - Running and managing Camel Running and deploying Camel Management and monitoring Part 6 - Out in the wild Clustering Microservices with Docker and Kubernetes Camel tooling Bonus online chapters Available at https://www.manning.com/books/camel-in-action-second-edition and in electronic versions of this book: Reactive Camel Camel and the IoT by Henryk Konsek

This timely volume provides a review of the state-of-the-art frameworks and methodologies for connecting diverse objects and devices according to the vision for an Internet of Things (IoT). A specific focus is placed on the communication, security, and privacy aspects of device connectivity in distributed environments. Insights and case studies are provided by an authoritative selection of contributors of international repute into the latest research advances and practical approaches with respect to the connectivity of heterogeneous smart and sensory devices. Topics and features: Examines aspects of device connectivity within the IoT Presents a resource-based architecture for IoT, and proposes a resource management framework for corporate device clouds Reviews integration approaches for the IoT environment, and discusses performance optimization of intelligent home networks Introduces a novel solution for interoperable data management in multi-clouds, and suggests an approach that addresses the debate over network neutrality in the IoT Describes issues of data security, privacy, access control, and authentication in the distributed IoT environment Reviews the evolution of VANETs in relation to the Internet of Vehicles, and provides a perspective on developing smart sustainable cities This invaluable text/reference will be of great benefit to a broad audience, from students and researchers interested in the IoT vision, to practicing communication engineers and network security specialists.

The book contains 10 chapters, and it is divided into four sections. The first section includes three chapters, providing an overview of Energy Management of Distributed Systems. It outlines typical concepts, such as Demand-Side Management, Demand Response, Distributed, and Hierarchical Control for Smart Micro-Grids. The second section contains three chapters and presents different control algorithms, software architectures, and simulation tools dedicated to Energy Management Systems. In the third section, the importance and the role of energy storage technology in a Distribution System, describing and comparing different types of energy storage systems, is shown. The fourth section shows how to identify and address potential threats for a Home Energy Management System. Finally, the fifth section discusses about Economical Optimization of Operational Cost for Micro-Grids, pointing out the effect of renewable energy sources, active loads, and energy storage systems on economic operation.

Preparing for the Certified OpenStack Administrator Exam

RabbitMQ Cookbook

Set Up and Manage Your OpenStack Cloud

The Internet of Things from a Distributed Computing Perspective

RabbitMQ in Action

Red Hat RHCSA 8 Cert Guide

For many organizations, a big part of DevOps' appeal is software automation using infrastructure-as-code techniques. This book presents developers, architects, and infra-ops engineers with a more practical option. You'll learn how a container-centric approach from OpenShift, Red Hat's cloud-based PaaS, can help your team deliver quality software through a self-service view of IT infrastructure. Three OpenShift experts at Red Hat explain how to configure Docker application containers and the Kubernetes cluster manager with OpenShift's developer- and operational-centric tools. Discover how this infrastructure-agnostic container management platform can help companies navigate the murky area where infrastructure-as-code ends and application automation begins. Get an application-centric view of automation—and understand why it's important Learn patterns and practical examples for managing continuous deployments such as rolling, A/B, blue-green, and canary Implement continuous integration pipelines with OpenShift's Jenkins capability Explore mechanisms for separating and managing configuration from static runtime software Learn how to use and customize OpenShift's source-to-image capability Delve into management and operational considerations when working with OpenShift-based application workloads Install a self-contained local version of the OpenShift environment on your computer Design, deploy, and maintain your own private or public Infrastructure as a Service (IaaS), using the open source OpenStack platform. In this practical guide, experienced developers and OpenStack contributors show you how to build clouds based on reference architectures, as well as how to perform daily administration tasks. Designed for horizontal scalability, OpenStack lets you build a cloud by integrating several technologies. This approach provides flexibility, but knowing which options to use can be bewildering. Once you complete this book, you'll know the right questions to ask while you organize compute, storage, and networking resources. If you already know how to manage multiple Ubuntu machines and maintain MySQL, you're ready to: Set up automated deployment and configuration Design a single-node cloud controller Use metrics to improve scalability Explore compute nodes, network design, and storage Install OpenStack packages Use an example architecture to help simplify decision-making Build a working environment to explore an IaaS cloud Manage users, projects, and quotas Tackle maintenance, debugging, and network troubleshooting Monitor, log, backup, and restore

Build your own enterprise applications and integration flows with JBoss and its products About This Book Build fast, smart, and flexible applications using JBoss Couple one or more JBoss products to effectively solve various business problems Explore the JBoss product ecosystem for improving the performance of your projects Who This Book Is For If you are a Java developer who wants to have a complete view of the JBoss ecosystem or quickly explore a specific JBoss Product, then this is the book you want. Integrators and consultants, familiar with JBoss, who want integrate several JBoss products within their ongoing project will also find this book useful. What You Will Learn Create new applications or integrate existing systems with JBoss products Setup and manage a JBoss domain Setup and manage a JBoss Fuse cluster with Fabric and Apache Karaf Create and deploy OSGi applications on JBoss Fuse containersv Manage enterprise data with JBoss Datagrid Aggregate various data sources with JBoss Data virtualization to offer data as a service Optimize your business and workflows with both JBoss Business RulesManagement System and JBoss Business Process Management platforms. In Detail Have you often wondered what is the best JBoss product to solve a specific problem? Do you want to get started with a specific JBoss product and know how to integrate different JBoss products in your IT Systems? Then this is the book for you. Through hands-on examples from the business world, this guide presents details on the major products and how you can build your own Enterprise services around the JBoss ecosystem. Starting with an introduction to the JBoss ecosystem, you will gradually move on to developing and deploying clustered application on JBoss Application Server, and setting up high availability using undertow or HA proxy loadbalancers. As you are moving to a micro service architecture, you will be taught how to package existing Java EE applications as micro service using Swarm or create your new micro services from scratch by coupling most popular Java EE frameworks like JPA, CDI with Undertow handlers. Next, you will install and configure JBoss Data grid in development and production environments, develop cache based applications and aggregate various data source in JBoss data virtualization. You will learn to build, deploy, and monitor integration scenarios using JBoss Fuse and run both producers/consumers applications relying on JBoss AMQ. Finally, you will learn to develop and run business workflows and make better decisions in your applications using Drools and Jboss BPM Suite Platform. Style and Approach The book works through the major JBoss products, with examples and instructions to help you understand each product and how they work together.

"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 provides a single customizable environment for building and 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 print book includes a free eBook in PDF, Kindle, and ePub formats 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 Airflow perfect for any data management task. About the book 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, and how to customize them for your pipeline's needs. 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 authors Bas Harenslak and Julian de Ruiter are data engineers with 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 Communicating with external systems 8 Building custom 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

WebSocket

Exam 210-451 and Exam 210-455

How Open Source Ate Software

Cloud Deployments Made Easy

Distributed Messaging for Everyone

Mastering MongoDB 3.x

Learn how to work with the Automate feature of CloudForms, the powerful Red Hat cloud management platform that lets you administer your virtual infrastructure, including hybrid public and private clouds. This practical hands-on introduction shows you how to increase your operational efficiency by automating day-to-day tasks that now require manual input. Throughout the book, author Peter McGowan provides a combination of theoretical information and practical coding examples to help you learn the Automate object model. With this CloudForms feature, you can create auto-scalable cloud applications, eliminate manual decisions and operations when provisioning virtual machines and cloud instances, and manage your complete virtual machine lifecycle. In six parts, this book helps you: Learn the objects and concepts for developing automation scripts with CloudForms Automate Customize the steps and workflows involved in provisioning virtual machines Create and use service catalogs, items, dialogs, objects, bundles, and hierarchies Use CloudForm's updated workflow to retire and delete virtual machines and services Orchestrate and coordinate with external services as part of a workflow Explore distributed automation processing as well as argument passing and handling Microsoft Azure Essentials from Microsoft Press is a series of free ebooks designed to help you advance your technical skills with Microsoft Azure. The first ebook in the series, Microsoft Azure Essentials: Fundamentals of Azure, introduces developers and IT professionals to the wide range of capabilities in Azure. The authors - both Microsoft MVPs in Azure - present both conceptual and how-to content for key areas, including: Azure Websites and Azure Cloud Services Azure Virtual Machines Azure Storage Azure Virtual Networks Databases Azure Active Directory Management tools Business scenarios Watch Microsoft Press's blog and Twitter (@MicrosoftPress) to learn about other free ebooks in the "Microsoft Azure Essentials" series.

IBM® Power Virtualization Center (IBM® PowerVCTM) is an advanced enterprise virtualization management offering for IBM Power Systems. This IBM Redbooks® publication introduces IBM PowerVC and helps you understand its functions, planning, installation, and setup. It also shows how IBM PowerVC can integrate with systems management tools such as Ansible or Terraform and that it also integrates well into a OpenShift container environment. IBM PowerVC Version 2.0.0 supports both large and small deployments, either by managing IBM PowerVM® that is controlled by the Hardware Management Console (HMC), or by IBM PowerVM NovaLink. With this capability, IBM PowerVC can manage IBM AIX®, IBM i, and Linux workloads that run on IBM POWER® hardware. IBM PowerVC is available as a Standard Edition, or as a Private Cloud Edition. IBM PowerVC includes the following features and benefits: Virtual image capture, import, export, deployment, and management Policy-based virtual machine (VM) placement to improve server usage Snapshots and cloning of VMs or volumes for backup or testing purposes Support of advanced storage capabilities such as IBM SVC vdisk mirroring of IBM Global Mirror Management of real-time optimization and VM resilience to increase productivity VM Mobility with placement policies to reduce the burden on IT staff in a simple-to-install and easy-to-use graphical user interface (GUI) Automated Simplified Remote Restart for improved availability of VMs ifor when a host is down Role-based security policies to ensure a secure environment for common tasks The ability to enable an administrator to enable Dynamic Resource Optimization on a schedule IBM PowerVC Private Cloud Edition includes all of the IBM PowerVC Standard Edition features and enhancements: A self-service portal that allows the provisioning of new VMs without direct system administrator intervention. There is an option for policy approvals for the requests that are received from the self-service portal. Pre-built deploy templates that are set up by the cloud administrator that simplify the deployment of VMs by the cloud user. Cloud management policies that simplify management of cloud deployments. Metering data that can be used for chargeback. This publication is for experienced users of IBM PowerVM and other virtualization solutions who want to understand and implement the next generation of enterprise virtualization management for Power Systems. Unless stated otherwise, the content of this publication refers to IBM PowerVC Version 2.0.0.

This IBM® Redpaper Redbooks® publication presents the IBM PowerKVM virtualization for scale-out Linux systems, including the new LC IBM Power Systems™. PowerKVM is open source server virtualization that is based on the IBM POWER8® processor technology. It includes the Linux open source technology of KVM virtualization, and it complements the performance, scalability, and security qualities of Linux. This book describes the concepts of PowerKVM and how you can deploy your virtual machines with the software stack included in the product. It helps you install and configure PowerKVM on your Power Systems server and provides guidance for managing the supported virtualization features by using the web interface and command-line interface (CLI). This information is for professionals who want to acquire a better understanding of PowerKVM virtualization technology to optimize Linux workload consolidation and use the POWER8 processor features. The intended audience also includes people in these roles: Clients Sales and marketing professionals Technical support professionals IBM Business Partners Independent software vendors Open source community IBM OpenPower partners It does not replace the latest marketing materials and configuration tools. It is intended as an additional source of information that, along with existing sources, can be used to increase your knowledge of IBM virtualization solutions. Before you start reading, you must be familiar with the general concepts of kernel-based virtual machine (KVM), Linux, and IBM Power architecture.

Spring Integration in Action

RabbitMQ in Depth

Understand the Open Source Movement and So Much More

It Professionals' Handbook

IBM PowerKVM: Configuration and Use

Master the objectives required to pass the Certified OpenStack Administrator exam. About This Book Focuses on providing a clear, concise strategy so you gain the specific skills required to pass the Certified OpenStack Administrator exam Includes exercises and performance-based tasks to ensure all exam objectives can be completed via the Horizon dashboard and command-line interface Includes a free OpenStack Virtual Appliance to practice the objectives covered throughout the book Includes a practice exam to put your OpenStack skills to the test to prove you have what it takes to conquer the live exam Updated for the 2017 exam featuring OpenStack Newton Who This Book Is For This book is for IT professionals, system administrators, DevOps engineers, and software developers with basic Linux command-line and networking knowledge. It's also a great guide for those interested in an entry-level OpenStack position but have limited real-world OpenStack experience. After passing the exam, Certified OpenStack Administrators will prove they have the required skills for the job. What You Will Learn Manage the Keystone identity service by creating and modifying domains, groups, projects, users, roles, services, endpoints, and quotas. Upload Glance images, launch new Nova instances, and create flavors, key pairs, and snapshots. Discover Neutron tenant and provider networks, security groups, routers, and floating IPs. Manage the Cinder block storage service by creating volumes and attaching them to instances. Create Swift containers and set access control lists to allow read/write access to your objects. Explore Heat orchestration templates and create, list, and update stacks. In Detail This book provides you with a specific strategy to pass the OpenStack Foundation's first professional certification: the Certified OpenStack Administrator. In a recent survey, 78% of respondents said the OpenStack skills shortage had deterred them from adopting OpenStack. Consider this an opportunity to increase employer and customer confidence by proving you have the skills required to administrate real-world OpenStack clouds. You will begin your journey by getting well-versed with the OpenStack environment, understanding the benefits of taking the exam, and installing an included OpenStack all-in-one virtual appliance so you can work through objectives covered throughout the book. After exploring the basics of the individual services, you will be introduced to strategies to accomplish the exam objectives relevant to Keystone, Glance, Nova, Neutron, Cinder, Swift, Heat, and troubleshooting. Finally, you'll benefit from the special tips section and a practice exam to put your knowledge to the test. By the end of the journey, you will be ready to become a Certified OpenStack Administrator! Style and approach Clear, concise, and straightforward with supporting diagrams and lab environment tutorials, this book will help you confidently pass Certified OpenStack Administrator objectives on the Horizon dashboard and command-line interface.

Until recently, creating desktop-like applications in the browser meant using inefficient Ajax or Comet technologies to communicate with the server. With this practical guide, you'll learn how to use WebSocket, a protocol that enables the client and server to communicate with each other on a single connection simultaneously. No more asynchronous communication or long polling! For developers with a good grasp of JavaScript (and perhaps Node.js), author Andrew Lombardi provides useful hands-on examples throughout the book to help you get up to speed with the WebSocket API. You'll also learn how to use WebSocket with Transport Layer Security (TLS). Learn how to use WebSocket API events, messages, attributes, and methods within your client application Build bi-directional chat applications on the client and server with WebSocket as the communication layer Create a subprotocol over WebSocket for STOMP 1.0, the Simple Text Oriented Messaging Protocol Use options for older browsers that don't natively support WebSocket Protect your WebSocket application against various attack vectors with TLS and other tools Debug applications by learning aspects of the WebSocket lifecycle

An expert's guide to build fault tolerant MongoDB application About This Book Master the advanced modeling, querying, and administration techniques in MongoDB and become a MongoDB expert Covers the latest updates and Big Data features frequently used by professional MongoDB developers and administrators If your goal is to become a certified MongoDB professional, this book is your perfect companion Who This Book Is For Mastering MongoDB is a book for database developers, architects, and administrators who want to learn how to use MongoDB more effectively and productively. If you have experience in, and are interested in working with, NoSQL databases to build apps and websites, then this book is for you. What You Will Learn Get hands-on with advanced querying techniques such as indexing, expressions, arrays, and more. Configure, monitor, and maintain highly scalable MongoDB environment like an expert. Master replication and data sharding to optimize read/write performance. Design secure and robust applications based on MongoDB. Administer MongoDB-based applications on-premise or in the cloud Scale MongoDB to achieve your design goals Integrate MongoDB with big data sources to process huge amounts of data In Detail MongoDB has grown to become the de facto NoSQL database with millions of users—from small startups to Fortune 500 companies. Addressing the limitations of SQL schema-based databases, MongoDB pioneered a shift of focus for DevOps and offered sharding and replication maintainable by DevOps teams. The book is based on MongoDB 3.x and covers topics ranging from database querying using the shell, built in drivers, and popular ODM mappers to more advanced topics such as sharding, high availability, and integration with big data sources. You will get an overview of MongoDB and how to play to its strengths, with relevant use cases. After that, you will learn how to query MongoDB effectively and make use of indexes as much as possible. The next part deals with the administration of MongoDB installations on-premise or in the cloud. We deal with database internals in the next section, explaining storage systems and how they can affect performance. The last section of this book deals with replication and MongoDB scaling, along with integration with heterogeneous data sources. By the end this book, you will be equipped with all the required industry skills and knowledge to become a certified MongoDB developer and administrator. Style and approach This book takes a practical, step-by-step approach to explain the concepts of MongoDB. Practical use-cases involving real-world examples are used throughout the book to clearly explain theoretical concepts.

Summary RabbitMQ in Depth is a practical guide to building and maintaining message-based applications. This book provides detailed coverage of RabbitMQ with an emphasis on why it works the way it does. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About

the Technology At the heart of most modern distributed applications is a queue that buffers, prioritizes, and routes message traffic. RabbitMQ is a high-performance message broker based on the Advanced Message Queuing Protocol. It's battle tested, ultrafast, and powerful enough to handle anything you can throw at it. It requires a few simple setup steps, and you can instantly start using it to manage low-level service communication, application integration, and distributed system message routing. About the Book RabbitMQ in Depth is a practical guide to building and maintaining message-based applications. This book provides detailed coverage of RabbitMQ with an emphasis on why it works the way it does. You'll find examples and detailed explanations based in real-world systems ranging from simple networked services to complex distributed designs. You'll also find the insights you need to make core architectural choices and develop procedures for effective operational management. What's Inside AMQP, the Advanced Message Queuing Protocol Communicating via MQTT, Stomp, and HTTP Valuable troubleshooting techniques Database integration About the Reader Written for programmers with a basic understanding of messaging-oriented systems. About the Author Gavin M. Roy is an active, open source evangelist and advocate who has been working with internet and enterprise technologies since the mid-90s. Technical editor James Titcumb is a freelance developer, trainer, speaker, and active contributor to open source projects. Table of Contents PART 1 – RABBITMQ AND APPLICATION ARCHITECTURE Foundational RabbitMQ How to speak Rabbit: the AMQ Protocol An in-depth tour of message properties Performance trade-offs in publishing Don't get messages; consume them Message patterns via exchange routing PART 2 – MANAGING RABBITMQ IN THE DATA CENTER OR THE CLOUD Scaling RabbitMQ with clusters Cross-cluster message distribution PART 3 – INTEGRATIONS AND CUSTOMIZATION Using alternative protocols Database integrations

Designing, Building, and Deploying Messaging Solutions

Ansible: Up and Running

DevOps with OpenShift

CCNA Cloud Complete Study Guide

Lightweight Client-Server Communications

Mastering CloudForms Automation

Pro Python System Administration, Second Edition explains and shows how to apply Python scripting in practice. It will show you how to approach and resolve real-world issues that most system administrators will come across in their careers. This book has been updated using Python 2.7 and Python 3 where appropriate. It also uses various new and relevant open source projects and tools that should now be used in practice. In this updated edition, you will find several projects in the categories of network administration, web server administration, and monitoring and database management. In each project, the author will define the problem, design the solution, and go through the more interesting implementation steps. Each project is accompanied by the source code of a fully working prototype, which you ' ll be able to use immediately or adapt to your requirements and environment. This book is primarily aimed at experienced system administrators whose day-to-day tasks involve looking after and managing small-to-medium-sized server estates. It will also be beneficial for system administrators who want to learn more about automation and want to apply their Python knowledge to solve various system administration problems. Python developers will also benefit from reading this book, especially if they are involved in developing automation and management tools.

A practical book filled with advanced recipes as well as plenty of code and real-life examples which will make your learning curve quick and easy. If you are a software developer who wants to develop distributed applications based on messaging [BISAC]; then this book is for you. It ' s assumed that you have some experience with multithreading applications and distributed applications. You are also expected to know the basic concepts of Web and cloud applications in order to follow the recipes effectively.

Reactive systems and event-driven architecture are becoming indispensable to application design, and companies are taking note. Reactive systems ensure that applications are responsive, resilient, and elastic no matter what failures or errors may be occurring, while event-driven architecture offers a flexible and composable option for distributed systems. This practical book helps Java developers bring these approaches together using Quarkus 2.x, the Kubernetes-native Java framework. Clement Escoffier and Ken Finnigan show you how to take advantage of event-driven and reactive principles to build robust distributed systems, reducing latency and increasing throughput, particularly in microservices and serverless applications. You'll also get a foundation in Quarkus to help you create true Kubernetes-native applications for the cloud. Understand the fundamentals of reactive systems and event-driven architecture Learn how to use Quarkus to build reactive applications Combine Quarkus with Apache Kafka or AMQP to build reactive systems Develop microservices that utilize messages with Quarkus for use in event-driven architectures Learn how to integrate external messaging systems, such as Apache Kafka, with Quarkus Build applications with Quarkus using reactive systems and reactive programming concepts

Learn how free software became open source and how you can sell open source software. This book provides a historical context of how open source has thoroughly transformed how we write software, how we cooperate, how we communicate, how we organize, and, ultimately, how we think about business values. You ' ll look at project and community examples including Linux, BSD, Apache, and Kubernetes, understand the open source development model, and how open source has influenced approaches more broadly, even proprietary software, such as open betas. You'll also examine the flipside, the "Second Machine Age," and the challenges of open source-based business models. Today, open source serves as shorthand for much broader trends and behaviors. It ' s not just about a free (in all senses of the word) alternative to commercial software. It increasingly is the new commercial software. How Open Source Ate Software reveals how open source has much in common, and is often closely allied, with many other trends in business and society. You'll see how it enables projects that go beyond any individual company. That makes open source not just a story about software, but a story about almost everything. What You'll Learn Understand open source opportunities and challenges Sell software if you ' re giving it away Apply open source principles more broadly to openorg, devops, etc. Review which organizational incentives you can implement Who This Book Is For Anyone who has an interest in what is happening in open source and the open source community, and anyone who is contemplating making a business that involves open source.

Learning MCollective

An Essential Guide for Cloud Administrators

Camel in Action

The Necropsy Book

Pro Python System Administration

Automating Configuration Management and Deployment the Easy Way

Teaches you how and what to study in order to be best prepared for the Certified OpenStack Administrator exam. This fast-growing technology is creating a market that needs more qualified IT specialists with proven skills. This book covers 100% of the exam requirements for both The OpenStack Foundation and the Mirantis OpenStack Certification Exam. Each theme is taught using practical exercises and instructions for the command line and for the graphical client (Horizon). Each chapter is followed by review questions, complete with answers. Even after you have taken and passed your OpenStack exam, this book will remain a useful reference. What You Will Learn Understand the components that make up the cloud. Install and make an OpenStack distribution from Mirantis, Red Hat or another community version. Work with OpenStack Identity Management, Dashboard, CLI, Object Storage, Block Storage, Networking, Telemetry, Orchestration, and Image Services. Learn how to troubleshoot all the main OpenStack services. Understand where to find information for future work with OpenStack. Who This Book Is For Certified OpenStack Administrator Study Guide is for Cloud and Linux engineers looking for a better understanding of how to work with the modern OpenStack IaaS Cloud, and wants to prove their knowledge by passing a Certified OpenStack Administrator Exam.

Learn Chef Provisioning like a boss and discover how to deploy software and manage hosts, along with engaging recipes to automate your cloud and server infrastructure with Chef. About This Book Leverage the power of Chef to transform your infrastructure into code to deploy new features in minutes Get step-by-step instructions to configure, deploy, and scale your applications Master specific Chef techniques to run an entire fleet of machines without breaking a sweat. Who This Book Is For If you are a system administrator, Linux administrator, a cloud developer, or someone who just wants to learn and apply Chef automation to your existing or new infrastructure, then this learning path will show you all you need to know. In order to get the most out of this learning path, some experience of programming or scripting languages would be useful. What You Will Learn Install Chef server on your own hosts Integrate Chef with cloud services Debug your cookbooks and Chef runs using the numerous inspection and logging facilities of Chef Extend Chef to meet your advanced needs by creating custom plugins for Knife and Ohai Create a perfect model system Use the best test-driven development methodologies In Detail Chef is a configuration management tool that turns IT infrastructure into code. Chef provides tools to manage systems at scale. This learning path takes you on a comprehensive tour of Chef's functionality, ranging from its core features to advanced development. You will be brought up to speed with what's new in Chef and how to set up your own Chef infrastructure for individuals, or small or large teams. You will learn to use the basic Chef command-line tools. We will also take you through the core concepts of managing users, applications, and your entire cloud infrastructure. You will learn the techniques of the pros by walking you through a host of step-by-step guides to solve real-world infrastructure automation challenges. You will learn to automate and document every aspect of your network, from the hardware to software, middleware, and all your containers. You will become familiar with the Chef's Provisioning tool. By the end of this course, you will be confident in how to manage your infrastructure, scale using the cloud, and extend the built-in functionality of Chef itself. The books used in this Learning Path are: 1) Chef Essentials 2) Chef Infrastructure Automation Cookbook – Second Edition 3) Mastering Chef Provisioning Style and approach This fast-paced guide covers the many facets of Chef and will teach administrators to use Chef as a birds-eye lens for their entire system. This book takes you through a host of step-by-step guides to solve real-world infrastructure automation challenges and offers elegant, time-saving solutions for a perfectly described and automated network.

Summary Spring Integration in Action is a hands-on guide to Spring-based messaging and integration. After addressing the core messaging patterns, such as those used in transformation and routing, the book turns to the adapters that enable integration with external systems. Readers will explore real-world enterprise integration scenarios using JMS, Web Services, file systems, and email. They will also learn about Spring Integration's support for working with XML. The book concludes with a practical guide to advanced topics such as concurrency, performance, system-management, and monitoring. The book features a foreword by Rod Johnson, Founder of the Spring Network. About the Technology Spring Integration extends the Spring Framework to support the patterns described in Gregor Hohpe and Bobby Woolf's Enterprise Integration Patterns. Like the Spring Framework itself, it focuses on developer productivity, making it easier to build, test, and maintain enterprise integration solutions. About the Book Spring Integration in Action is an introduction and guide to enterprise integration and messaging using the Spring Integration framework. The book starts off by reviewing core messaging patterns, such as those used in transformation and routing. It then drills down into real-world enterprise integration scenarios using JMS, Web Services, filesystems, email, and more. You'll find an emphasis on testing, along with practical coverage of topics like concurrency, scheduling, system management, and monitoring. This book is accessible to developers who know Java. Experience with Spring and EIP is helpful but not assumed. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Realistic examples Expert advice from Spring Integration creators Detailed coverage of Spring Integration 2 features About the Authors Mark Fisher is the Spring Integration founder and project lead. Jonas Partner, Marius Bogoevici, and Iwein Fuld have all been project committers and are recognized experts on Spring and Spring Integration. Table of Contents PART 1 BACKGROUND Introduction to Spring Integration Enterprise integration fundamentals 24 PART 2 MESSAGING Messages and channels Message Endpoints Getting down to business Go beyond sequential processing: routing and filtering Splitting and aggregating messages PART 3 INTEGRATING SYSTEMS Handling messages with XML payloads Spring Integration and the Java Message Service Email-based integration Filesystem integration Spring Integration and web services Chatting and tweeting PART 4 ADVANCED TOPICS Monitoring and management Managing scheduling and concurrency Batch applications and enterprise integration Scaling messaging applications with OSGi Testing

Today, organizations are responding to market demands and regulatory requirements faster than ever by extending their applications and data to new digital applications. This drive to deliver new functions at speed has paved the way for a huge growth in cloud-native applications, hosted in both public and private cloud infrastructures. Leading organizations are now exploiting the best of both worlds by combining their traditional enterprise IT with cloud. This hybrid cloud approach places new requirements on the integration architectures needed to bring these two worlds together. One of the largest providers of application logic and data services in enterprises today is IBM Z, making it a critical service provider in a hybrid cloud architecture. The primary goal of this IBM Redpaper publication is to help IT architects choose between the different application integration architectures that can be used for hybrid integration with IBM Z, including REST APIs, messaging, and event streams.

Mastering RabbitMQ

Certified OpenStack Administrator Study Guide

Accelerating Modernization with Agile Integration

An expert's guide to building fault-tolerant MongoDB applications

Chef: Powerful Infrastructure Automation

SAN/LAN Monthly Newsletter September 2010

Cisco has announced big changes to its certification program. As of February 24, 2020, all current certifications will be retired, and Cisco will begin offering new certification programs. The good news is if you're working toward any current CCNA certification, keep going. You have until February 24, 2020 to complete your current CCNA. If you already have CCENT/ICND1 certification and would like to earn CCNA, you have until February 23, 2020 to complete your CCNA certification in the current program. Likewise, if you're thinking of completing the current CCENT/ICND1, ICND2, or CCNA Routing and Switching certification, you can still complete them between now and February 23, 2020. Increase the value of your organization's cloud network—and invest in your education The Cisco Cloud certification validates the skill set of individuals on industry-leading cloud solutions and best practices, as well as offering job role-based curricula for all levels of an IT staff. CCNA Cloud Complete Study Guide prepares you to take two required exams: 210-451, Understanding Cisco Cloud Fundamentals, and 210-455, Introducing Cisco Cloud Administration. It covers everything you can expect to encounter on the exams and also gives you a year of FREE access to Sybex's superior online interactive learning environment and test bank, including chapter tests, practice exams, a glossary of key terms, and electronic flashcards. Cisco's CCNA Cloud certification covers cloud characteristics and models, cloud deployment, and basic knowledge of cloud compute, cloud networking, and cloud storage. It also covers cloud infrastructure administration and reporting, chargeback and billing reports, cloud provisioning, cloud systems management and monitoring, and cloud remediation. With thorough coverage, practical instruction, and expert insight, this book provides an ideal resource for Exam 210-451 and Exam 210-455 preparation. • Includes an opening list of exam topics • Provides valuable hands-on exercises • Offers practical real-world examples • Distills in-depth perspective from cloud computing experts This book is the perfect resource for anyone seeking to earn the challenging, but rewarding CCNA Cloud certification.

This is the eBook version of the print title. Learn, prepare, and practice for Red Hat RHCSA 8 (EX200) exam success with this Cert Guide from Pearson IT Certification, a leader in IT Certification learning. Master Red Hat RHCSA 8 EX200 exam topics Assess your knowledge with chapter-ending quizzes Review key concepts with exam-preparation tasks Practice with four unique practice tests Learn from two full hours of video training from the author's Red Hat Certified System Administrator (RHCSA) Complete Video Course, 3rd Edition. Red Hat RHCSA 8 Cert Guide is a best-of-breed exam study guide. Leading Linux consultant, author, and instructor Sander van Vugt shares preparation hints and test-taking tips, helping you identify areas of weakness and improve both your conceptual knowledge and hands-on skills. Material is presented in a concise manner, focusing on increasing your understanding and retention of exam topics. The book presents you with an organized test-preparation routine through the use of proven series elements and techniques. Exam topic lists make referencing easy. Chapter-ending Exam Preparation Tasks help you drill on key concepts you must know thoroughly. Review questions help you assess your knowledge, and a final preparation chapter guides you through tools and resources to help you craft your final study plan. Well regarded for its level of detail, assessment features, and challenging review questions and exercises, this study guide helps you master the concepts and techniques that will enable you to succeed on the exam the first time, including Basic system management: Installation, tools, file management, text files, RHEL8 connections, user/group management, permissions, and network configuration Operating running systems: Managing software, processes, storage, and advanced storage; working with systemd; scheduling tasks; and configuring logging Advanced system administration: Managing the kernel and boot procedures, essential troubleshooting, bash shell scripting Managing network services: Configuring SSH, firewalls, and time services; managing Apache HTTP services and SE Linux; and accessing network storage

Microservices is an architectural style in which large, complex software applications are composed of one or more smaller services. Each of these microservices focuses on completing one task that represents a small business capability. These microservices can be developed in any programming language. This IBM® Redbooks® publication covers Microservices best practices for Java. It focuses on creating cloud native applications using the latest version of IBM WebSphere® Application Server Liberty, IBM Bluemix® and other Open Source Frameworks in the Microservices ecosystem to highlight Microservices best practices for Java.

The organization pursuing digital transformation must embrace new ways to use and deploy integration technologies, so they can move quickly in a manner appropriate to the goals of multicloud, decentralization, and microservices. The integration layer must transform to allow organizations to move boldly in building new customer experiences, rather than forcing models for architecture and development that pull away from maximizing the organization's productivity. Many organizations have started embracing agile application techniques, such as microservice architecture, and are now seeing the benefits of that shift. This approach complements and accelerates an enterprise's API strategy. Businesses should also seek to use this approach to modernize their existing integration and messaging infrastructure to achieve more effective ways to manage and operate their integration services in their private or public cloud. This IBM® Redbooks® publication explores the merits of what we refer to as agile integration; a container-based, decentralized, and microservice-aligned approach for integration solutions that meets the demands of agility, scalability, and resilience required by digital transformation. It also discusses how the IBM Cloud Pak for Integration marks a significant leap forward in integration technology by embracing both a cloud-native approach and container technology to achieve the goals of agile integration. The target audiences for this book are cloud integration architects, IT specialists, and application developers.

Energy Management of Distributed Generation Systems

IBM PowerVC Version 2.0 Introduction and Configuration

IBM Cloud Private System Administrator's Guide

Microservices Best Practices for Java

Gigabit/ATM Monthly Newsletter September 2010

Consize Cloud Compute

Enterprise Integration Patterns provides an invaluable catalog of sixty-five patterns, with real-world solutions that demonstrate the formidable of messaging and help you to design effective messaging solutions for your enterprise. The authors also include examples covering a variety of different integration technologies, such as JMS, MSM, TIBCO ActiveEnterprise, Microsoft BizTalk, SOAP, and XSL. A case study describing a bond trading system illustrates the patterns in practice, and the book offers a look at emerging standards, as well as insights into what the future of enterprise integration might hold. This book provides a consistent vocabulary and visual notation framework

to describe large-scale integration solutions across many technologies. It also explores in detail the advantages and limitations of asynchronous messaging architectures. The authors present practical advice on designing code that connects an application to a messaging system, and provide extensive information to help you determine when to send a message, how to route it to the proper destination, and how to monitor the health of a messaging system. If you want to know how to manage, monitor, and maintain a messaging system once it is in use, get this book.

Mastering CloudForms AutomationAn Essential Guide for Cloud Administrators"O'Reilly Media, Inc."

Orchestrate change across server clusters in near realtime with MCollective, the framework that works in concert with Puppet, Chef, and other configuration management tools. Ideal for system administrators and operations or DevOps engineers at any level, this hands-on guide teaches you how to build and test a real installation of MCollective servers and clients in your environment. Learn how to build an entire installation by hand, know where every configuration file lives, and understand every configuration parameter and what it means. Whether you manage a small environment or one that's immense in scale, this book shows you how to orchestrate specific actions faster and better than you do now. Tour MCollective's architecture, backbone, transport, and security controls Configure MCollective components to match your production environment Create and use collectives to handle thousands of remote MCollective agents Use ActiveMQ Network of Brokers to resolve multi-site or redundancy requirements Learn how to use community-built client and agent plugins, with concrete examples Create your own server and client plugins to perform a variety of actions Learn recommended best practices for using MCollective

Every enterprise application creates data, whether it's log messages, metrics, user activity, outgoing messages, or something else. And how to move all of this data becomes nearly as important as the data itself. If you're an application architect, developer, or production engineer new to Apache Kafka, this practical guide shows you how to use this open source streaming platform to handle real-time data feeds. Engineers from Confluent and LinkedIn who are responsible for developing Kafka explain how to deploy production Kafka clusters, write reliable event-driven microservices, and build scalable stream-processing applications with this platform. Through detailed examples, you'll learn Kafka's design principles, reliability guarantees, key APIs, and architecture details, including the replication protocol, the controller, and the storage layer. Understand publish-subscribe messaging and how it fits in the big data ecosystem. Explore Kafka producers and consumers for writing and reading messages Understand Kafka patterns and use-case requirements to ensure reliable data delivery Get best practices for building data pipelines and applications with Kafka Manage Kafka in production, and learn to perform monitoring, tuning, and maintenance tasks Learn the most critical metrics among Kafka's operational measurements Explore how Kafka's stream delivery capabilities make it a perfect source for stream processing systems

Reactive Systems in Java

Connectivity Frameworks for Smart Devices

OpenStack Operations Guide

IBM Z Integration Guide for Hybrid Cloud

Microsoft Azure Essentials - Fundamentals of Azure

A Guide for Veterinary Students, Residents, Clinicians, Pathologists, and Biological Researchers

Master the art of developing message-based applications with RabbitMQ About This Book Learn how to administer, manage, and extend your own message broker, RabbitMQ Develop clients to make a message bridge between your software systems using RabbitMQ Discover how to achieve proficiency with RabbitMQ with the well-

defined descriptions of the topics Who This Book Is For If you are an intermediate-level RabbitMQ developer, who wants to achieve professional-level expertise in the subject, this book is for you. You'll also need to have a decent understanding of message queuing. What You Will Learn Administer RabbitMQ using

different tools Understand the roots and details of messaging, message brokers, and AMQP protocol Scale the RabbitMQ server using the clusters and high availability techniques Extend RabbitMQ by developing the Erlang OTP-based applications that use the RabbitMQ API Manage the RabbitMQ server using its powerful tools

Monitor the RabbitMQ Server using different open source tools such as Nagios, Munin, and Zabbix Ensure your RabbitMQ's security using SSL, SASL, and access control Develop RabbitMQ clients using Java, Python, and C# with an industry example In Detail RabbitMQ is one of the most powerful Open Source message broker

software, which is widely used in tech companies such as Mozilla, VMware, Google, AT&T, and so on. RabbitMQ gives you lots of fantastic and easy-to-manage functionalities to control and manage the messaging facility with lots of community support. As scalability is one of our major modern problems, messaging with

RabbitMQ is the main part of the solution to this problem. This book explains and demonstrates the RabbitMQ server in a detailed way. It provides you with lots of real-world examples and advanced solutions to tackle the scalability issues. You'll begin your journey with the installation and configuration of the

RabbitMQ server, while also being given specific details pertaining to the subject. Next, you'll study the major problems that our server faces, including scalability and high availability, and try to get the solutions for both of these issues by using the RabbitMQ mechanisms. Following on from this, you'll get to

design and develop your own plugins using the Erlang language and RabbitMQ's internal API. This knowledge will help you to start with the management and monitoring of the messages, tools, and applications. You'll also gain an understanding of the security and integrity of the messaging facilities that RabbitMQ

provides. In the last few chapters, you will build and keep track of your clients (senders and receivers) using Java, Python, and C#. Style and approach An easy-to-follow guide, full of hands-on examples based around managing, monitoring, extending, and securing RabbitMQ and its internal tools. You will learn how to

develop your own clients using Java, Python, and C#.

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

interoperability challenges. About the Technology There's a virtual switchboard at the core of most large applications where messages race between servers, programs, and services. RabbitMQ is an efficient and easy-to-deploy queue that handles this message traffic effortlessly in all situations, from web startups to

massive enterprise systems. About the Book RabbitMQ in Action teaches you to build and manage scalable applications in multiple languages using the RabbitMQ messaging server. It's a snap to get started. You'll learn how message queuing works and how RabbitMQ fits in. Then, you'll explore practical scalability and

interoperability issues through many examples. By the end, you'll know how to make Rabbit run like a well-oiled machine in a 24 x 7 x 365 environment. Written for developers familiar with Python, PHP, Java, .NET, or any other modern programming language. No RabbitMQ experience required. Purchase of the print book

comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Learn fundamental messaging design patterns Use patterns for on-demand scalability Glue a PHP frontend to a backend written in anything Implement a PubSub-alerting service in 30 minutes flat

Configure RabbitMQ's built-in clustering Monitor, manage, extend, and tune RabbitMQ =====?===== Table of Contents Pulling RabbitMQ out of the hat Understanding messaging Running and administering Rabbit Solving problems with Rabbit: coding and patterns Clustering and dealing with

failure Writing code that survives failure Warrens and Shovels: failover and replication Administering RabbitMQ from the Web Controlling Rabbit with the REST API Monitoring: Houston, we have a problem Supercharging and securing your Rabbit Smart Rabbits: extending RabbitMQ

IBM® Cloud Private is an application platform for developing and managing containerized applications across hybrid cloud environments, on-premises and public clouds. It is an integrated environment for managing containers that includes the container orchestrator Kubernetes, a private image registry, a management

console, and monitoring frameworks. This IBM Redbooks covers tasks performed by IBM Cloud Private system administrators such as installation for high availability, configuration, backup and restore, using persistent volumes, networking, security, logging and monitoring. Istio integration, troubleshooting and so on.

As part of this project we also developed several code examples and you can download those from the IBM Redbooks GitHub location: <https://github.com/IBMRedbooks>. The authors team has many years of experience in implementing IBM Cloud Private and other cloud solutions in production environments, so throughout this

document we took the approach of providing you the recommended practices in those areas. If you are an IBM Cloud Private system administrator, this book is for you. If you are developing applications on IBM Cloud Private, you can see the IBM Redbooks publication IBM Cloud Private Application Developer's Guide,

SG24-8441.

Optimized for Kubernetes, Quarkus is designed to help you create Java applications that are cloud first, container native, and serverless capable. With this cookbook, authors Alex Soto Bueno and Jason Porter from Red Hat provide detailed solutions for installing, interacting with, and using Quarkus in the development

and production of microservices. The recipes in this book show midlevel to senior developers familiar with Java enterprise application development how to get started with Quarkus quickly. You'll become familiar with how Quarkus works within the wider Java ecosystem and discover ways to adapt this framework to your

particular needs. You'll learn how to: Shorten the development cycle by enabling live reloading in dev mode Connect to and communicate with Kafka Develop with the reactive programming model Easily add fault tolerance to your services Build your application as a Kubernetes-ready container Ease development with OpenAPI

and test a native Quarkus application

Parallel Server Management in Puppet and Chef

Enterprise Integration Patterns

Data Pipelines with Apache Airflow

Kafka: The Definitive Guide

JBoss: Developer's Guide

Quarkus Cookbook

Among the many configuration management tools available, Ansible has some distinct advantages—it's minimal in nature, you don't need to install anything on your nodes, and it has an easy learning curve. This practical guide shows you how to be productive with this tool quickly, whether you're a developer deploying code to production or a system administrator looking for a better automation

solution. Author Lorin Hochstein shows you how to write playbooks (Ansible's configuration management scripts), manage remote servers, and explore the tool's real power: built-in declarative modules. You'll discover that Ansible has the functionality you need and the simplicity you desire. Understand how Ansible differs from other configuration management systems Use the YAML file format to

write your own playbooks Learn Ansible's support for variables and facts Work with a complete example to deploy a non-trivial application Use roles to simplify and reuse playbooks Make playbooks run faster with ssh multiplexing, pipelining, and parallelism Deploy applications to Amazon EC2 and other cloud platforms Use Ansible to create Docker images and deploy Docker containers

EX200

Real-Time Data and Stream Processing at Scale