

Openstack In Action

This book constitutes the refereed proceedings of the tracks and workshops which complemented the 14th European Conference on Software Architecture, ECSA 2020, held in L'Aquila, Italy*, in September 2020. The 30 full papers and 9 short papers presented in this volume were carefully reviewed and selected from 72 submissions. Papers presented were accepted into the following tracks and workshops: ECSA 2020 Doctoral Symposium track; ECSA 2020 Tool Demos track; ECSA 2020 Gender Diversity in Software Architecture & Software Engineering track; CASA - 3rd International Workshop on Context-aware, Autonomous and Smart Architecture; CSE/QUDOS - Joint Workshop on Continuous Software Engineering and Quality-Aware DevOps; DETECT - 3rd International Workshop on Modeling, Verification and Testing of Dependable Critical Systems; FAACS-MDE4SA - Joint Workshop on Formal Approaches for Advanced Computing Systems and Model-Driven Engineering for Software Architecture; IoT-ASAP - 4th International Workshop on Engineering IoT Systems: Architectures, Services, Applications, and Platforms; SASI4 - 2nd Workshop on Systems, Architectures, and Solutions for Industry 4.0; WASA - 6th International Workshop on Automotive System/Software Architecture. *The conference was held virtually due to the COVID-19 pandemic.

Over 90 practical, actionable recipes to automate, test, and manage your infrastructure quickly and effectively About This Book Bring down your delivery timeline from days to hours by treating your server configurations and VMs as code, just like you would with software code. Take your existing knowledge and skill set with your existing tools (Puppet, Chef, or Docker) to the next level and solve IT infrastructure challenges. Use practical recipes to use code to provision and deploy servers and applications and have greater control of your infrastructure. Who This Book Is For This book is for DevOps engineers and developers working in cross-functional teams or operations and would now switch to IAC to manage complex infrastructures. What You Will Learn Provision local and remote development environments with Vagrant Automate production infrastructures with Terraform, Ansible and Cloud-init on AWS, OpenStack, Google Cloud, Digital Ocean, and more Manage and test automated systems using Chef and Puppet Build, ship, and debug optimized Docker containers Explore the best practices to automate and test everything from cloud infrastructures to operating system configuration In Detail Infrastructure as Code (IAC) is a key aspect of the DevOps movement, and this book will show you how to transform the way you work with your infrastructure—by treating it as software. This book is dedicated to helping you discover the essentials of infrastructure automation and its related practices; the over 90 organized practical solutions will demonstrate how to work with some of the very best tools and cloud solutions. You will learn how to deploy repeatable infrastructures and services on AWS, OpenStack, Google Cloud, and Digital Ocean. You will see both Ansible and Terraform in action, manipulate the best bits from cloud-init to easily bootstrap instances, and simulate consistent environments locally or remotely using Vagrant. You will discover how to automate and test a range of system tasks using Chef or Puppet. You will also build, test, and debug various Docker containers having developers' interests in mind. This book will help you to use the right tools, techniques, and approaches to deliver working solutions for today's modern infrastructure challenges. Style and approach This is a recipe-based book that allows you to venture into some of the most cutting-edge practices and techniques about IAC and solve immediate problems when trying to implement them. 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

Undoubtedly, presupposition theory is a major chapter in the success story of dynamic semantics. This book features papers on this topic based on a conference on "Presupposition" convened in Stuttgart in October 2000.

Summary The best way to learn microservices development is to build something! Bootstrapping Microservices with Docker, Kubernetes, and Terraform guides you from zero through to a complete microservices project, including fast prototyping, development, and deployment. You'll get your feet wet using industry-standard tools as you learn and practice the practical skills you'll use for every microservices application. Following a true bootstrapping approach, you'll begin with a simple, familiar application and build up your knowledge and skills as you create and deploy a real microservices project. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Taking microservices from proof of concept to production is a complex, multi-step operation relying on tools like Docker, Terraform, and Kubernetes for packaging and deployment. The best way to learn the process is to build a project from the ground up, and that's exactly what you'll do with this book! About the book In Bootstrapping Microservices with Docker, Kubernetes, and Terraform, author Ashley Davis lays out a comprehensive approach to building microservices. You'll start with a simple design and work layer-by-layer until you've created your own video streaming application. As you go, you'll learn to configure cloud infrastructure with Terraform, package microservices using Docker, and deploy your finished project to a Kubernetes cluster. What's inside Developing and testing microservices applications Working with cloud providers Applying automated testing Implementing infrastructure as code and setting up a continuous delivery pipeline Monitoring, managing, and troubleshooting About the reader

Examples are in JavaScript. No experience with microservices, Kubernetes, Terraform, or Docker required. About the author Ashley Davis is a software developer, entrepreneur, stock trader, and the author of Manning's Data Wrangling with JavaScript. Table of Contents 1 Why microservices? 2 Creating your first microservice 3 Publishing your first microservice 4 Data management for microservices 5 Communication between microservices 6 Creating your production environment 7 Getting to continuous delivery 8 Automated testing for microservices 9 Exploring FlixTube 10 Healthy microservices 11 Pathways to scalability

OpenStack is today's leading technology for building and integrating public and private clouds. Common OpenStack Deployments is a complete, practical guide to deploying OpenStack and understanding its internals. Key project contributor Elizabeth Joseph, with expert implementer Matt Fischer, shares up-to-date recipes for deploying OpenStack on both virtual and physical servers, and for using OpenStack to address any real-world challenge. First, Joseph and Fischer help you master OpenStack concepts and components by guiding you through small-scale, virtualized deployments. As you deepen your understanding, they guide you through building large, horizontally scalable infrastructures that integrate multiple components in a feature-rich cloud environment. Throughout, you'll find up-to-the minute coverage of enhancements that make the OpenStack platform more mature and production ready, plus expert tips on debugging and growth. The authors conclude by introducing the broader OpenStack ecosystem, showing you how to drive value through hybrid clouds blending local and hosted solutions. Drawing on extensive personal experience, they address issues ranging from cost to data sovereignty and security. Common OpenStack Deployments is the ideal resource for all network and data center professionals who want to apply OpenStack in proof-of-concept or in production, and for every instructor or student who wants to leverage today's hottest cloud technology. Comprehensive coverage includes Customizing, deploying, and scaling OpenStack in any environment Quickly building single-server test deployments with DevStack Making the right networking decisions for any OpenStack deployment Manually installing Nova compute, Keystone identity, Glance image storage, and Neutron networking Setting up controllers and compute nodes Deploying private compute clouds with Puppet Metering clouds with Ceilometer Implementing block and object storage clouds Provisioning on bare metal with OpenStack Ironic Controlling containers with OpenStack Magnum Troubleshooting OpenStack: error messages, logs, tools, configuration problems, Puppet debugging, and more Step-by-step virtualized reference deployment using KVM/QEMU on Ubuntu

This book is written to help enterprise architects implement an OpenStack(r) cloud. With architects with one foot in information technology and the other in business operations in mind, we want to offer insights and best practices to help you achieve multiple (and sometimes competing) goals. If you're looking for vendor-neutral answers about planning your path to an OpenStack cloud, you're in the right place. Members of the OpenStack community-technologists, business leaders and product managers-collaborated on this book to explain how to get started with an OpenStack cloud. We've included pros and cons to help you make better choices when setting up your cloud, along with anticipated investments of both time and money. In this book, we'll discuss the considerations involved and how to make OpenStack cloud decisions about models, forming your team, organization and process changes, choosing workloads, and implementation from proof-of-concept through ongoing maintenance. Topics include: * Your technology options and their pros and cons * What to expect-in support, level of investment, and customization-from each type of cloud and consumption model * Operational models for a cloud, including staffing, plus how to manage consumption of cloud services in your business * How to assess the cloud's value to your business. After reading this book, you'll understand the process of building an OpenStack cloud, various cloud models, and operational and application approaches. You'll understand what decisions to make before building your cloud, and their effects on cost, resources and capabilities.

Even small applications have dozens of components. Large applications may have thousands, which makes them challenging to install, maintain, and remove. Docker bundles all application components into a package called a container that keeps things tidy and helps manage any dependencies on other applications or infrastructure. Docker in Action, Second Edition teaches you the skills and knowledge you need to create, deploy, and manage applications hosted in Docker containers. This bestseller has been fully updated with new examples, best practices, and entirely new chapters. You'll start with a clear explanation of the Docker model and learn how to package applications in containers, including techniques for testing and distributing applications. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Leverage the best SDN technologies for your OpenStack-based cloud infrastructure About This Book Learn how to leverage critical SDN technologies for OpenStack Networking APIs via plugins and drivers Champion the skills of achieving complete SDN with OpenStack with specific use cases and capabilities only covered in this title Discover exactly how you could implement cost-effective OpenStack SDN integration for your organization Who This Book Is For Administrators, and cloud operators who would like to implement Software Defined Networking on OpenStack clouds. Some prior experience of network infrastructure and networking concepts is assumed. What You Will Learn Understand how OVS is used for Overlay networks Get familiar with SDN Controllers with Architectural details and functionalities Create core ODL services and understand how OpenDaylight integrates with OpenStack to provide SDN capabilities Understand OpenContrail architecture and how it supports key SDN functionality such as Service Function Chaining (SFC) along with OpenStack Explore Open Network Operating System (ONOS) – a carrier grade SDN platform embraced by the biggest telecom service providers Learn about upcoming SDN technologies in OpenStack such as Dragonflow and OVN In Detail Networking is one the pillars of OpenStack and OpenStack Networking are designed to support programmability and Software-Defined Networks. OpenStack Networking has been evolving from simple APIs and functionality in Quantum to more complex capabilities in Neutron. Armed with the basic knowledge, this book will help the readers to explore popular SDN technologies, namely, OpenDaylight (ODL), OpenContrail, Open Network

Operating System (ONOS) and Open Virtual Network (OVN). The first couple of chapters will provide an overview of OpenStack Networking and SDN in general. Thereafter a set of chapters are devoted to OpenDaylight (ODL), OpenContrail and their integration with OpenStack Networking. The book then introduces you to Open Network Operating System (ONOS) which is fast becoming a carrier grade SDN platform. We will conclude the book with overview of upcoming SDN projects within OpenStack namely OVN and Dragonflow. By the end of the book, the readers will be familiar with SDN technologies and know how they can be leveraged in an OpenStack based cloud. Style and approach A hands-on practical tutorial through use cases and examples for Software Defined Networking with OpenStack.

OpenStack was created with the audacious goal of being the ubiquitous software choice for building public and private cloud infrastructures. In just over a year, it's become the most talked-about project in open source. This concise book introduces OpenStack's general design and primary software components in detail, and shows you how to start using it to build cloud infrastructures. If you're a developer, technologist, or system administrator familiar with cloud offerings such as Rackspace Cloud or Amazon Web Services, *Deploying OpenStack* shows you how to obtain and deploy OpenStack software in a few controlled scenarios. Learn about OpenStack Compute (known as "Nova"), OpenStack Object Store ("Swift"), and OpenStack Image Service ("Glance") Understand common pitfalls in architecting, deploying, and implementing your cloud infrastructure with OpenStack Determine which version of the OpenStack code base best suits your deployment needs Define your deployment scenario and finalize key design choices Install Nova on a single node with either the StackOps distro or an Ubuntu package Be familiar with important configuration options and important administrative commands

Keystone—OpenStack's Identity service—provides secure controlled access to a cloud's resources. In OpenStack environments, Keystone performs many vital functions, such as authenticating users and determining what resources users are authorized to access. Whether the cloud is private, public, or dedicated, access to cloud resources and security is essential. This practical guide to using Keystone provides detailed, step-by-step guidance to creating a secure cloud environment at the Infrastructure-as-a-Service layer—as well as key practices for safeguarding your cloud's ongoing security. Learn about Keystone's fundamental capabilities for providing Identity, Authentication, and Access Management Perform basic Keystone operations, using concrete examples and the latest version (v3) of Keystone's Identity API Understand Keystone's unique support for multiple token formats, including how it has evolved over time Get an in-depth explanation of Keystone's LDAP support and how to configure Keystone to integrate with LDAP Learn about one of Keystone's most sought-after features—support for federated identity

This book provides a complete and strategic overview of Multi-Access Edge Computing (MEC). It covers network and technology aspects, describes the market scenarios from the different stakeholders' point of view, and analyzes deployment aspects and actions to engage the ecosystem. MEC exists in and supports a highly complex "5G world" in which technologists and non-technology decision makers must act in concert and do so within a large interconnected ecosystem of which MEC is just one, albeit an important, part. Divided into three sections, with several chapters in each, the book addresses these three key aspects: technology, markets, and ecosystems.

Discover your complete guide to designing, deploying, and managing OpenStack-based clouds in mid-to-large IT infrastructures with best practices, expert understanding, and more About This Book Design and deploy an OpenStack-based cloud in your mid-to-large IT infrastructure using automation tools and best practices Keep yourself up-to-date with valuable insights into OpenStack components and new services in the latest OpenStack release Discover how the new features in the latest OpenStack release can help your enterprise and infrastructure Who This Book Is For This book is for system administrators, cloud engineers, and system architects who would like to deploy an OpenStack-based cloud in a mid-to-large IT infrastructure. This book requires a moderate level of system administration and familiarity with cloud concepts. What You Will Learn Explore the main architecture design of OpenStack components and core-by-core services, and how they work together Design different high availability scenarios and plan for a no-single-point-of-failure environment Set up a multinode environment in production using orchestration tools Boost OpenStack's performance with advanced configuration Delve into various hypervisors and container technology supported by OpenStack Get familiar with deployment methods and discover use cases in a real production environment Adopt the DevOps style of automation while deploying and operating in an OpenStack environment Monitor the cloud infrastructure and make decisions on maintenance and performance improvement In Detail In this second edition, you will get to grips with the latest features of OpenStack. Starting with an overview of the OpenStack architecture, you'll see how to adopt the DevOps style of automation while deploying and operating in an OpenStack environment. We'll show you how to create your own OpenStack private cloud. Then you'll learn about various hypervisors and container technology supported by OpenStack. You'll get an understanding about the segregation of compute nodes based on reliability and availability needs. We'll cover various storage types in OpenStack and advanced networking aspects such as SDN and NFV. Next, you'll understand the OpenStack infrastructure from a cloud user point of view. Moving on, you'll develop troubleshooting skills, and get a comprehensive understanding of services such as high availability and failover in OpenStack. Finally, you will gain experience of running a centralized logging server and monitoring OpenStack services. The book will show you how to carry out performance tuning based on OpenStack service logs. You will be able to master OpenStack benchmarking and performance tuning. By the end of the book, you'll be ready to take steps to deploy and manage an OpenStack cloud with the latest open source technologies. Style and approach This book will help you understand the flexibility of OpenStack by showcasing integration of several out-of-the-box solutions in order to build a large-scale cloud environment.. It will also cover detailed discussions on the various design and deployment strategies for implementing a fault-tolerant and highly available cloud infrastructure.

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.

This one-of-a-kind OpenStack self-assessment will make you the assured OpenStack domain veteran by revealing just what you need to know to be fluent and ready for any OpenStack challenge. How do I reduce the effort in the OpenStack work to be done to get problems solved? How can I ensure that plans of action include every OpenStack task and that every OpenStack outcome is in place? How will I save time investigating strategic and tactical options and ensuring OpenStack opportunity costs are low? How can I deliver tailored OpenStack advice instantly with structured going-forward plans? There's no better guide through these mind-expanding questions than acclaimed best-selling author Gerard Blokdyk. Blokdyk ensures all OpenStack essentials are covered, from every angle: the OpenStack self-assessment shows succinctly and clearly that what needs to be clarified to organize the business/project activities and processes so that OpenStack outcomes are achieved. Contains extensive criteria grounded in past and current successful projects and activities by experienced OpenStack practitioners. Their mastery, combined with the uncommon elegance of the self-assessment, provides its superior value to you in knowing how to ensure the outcome of any efforts in OpenStack are maximized with professional results. Your purchase includes access to the \$249 value OpenStack self-assessment dashboard download which gives you your dynamically prioritized projects-ready tool and shows your organization exactly what to do next. Your exclusive instant access details can be found in your book.

This book is a new contribution aiming to give some last research findings in the field of optimization and computing. This work is in the same field target than our two previous books published: Recent Developments in Metaheuristics and Metaheuristics for Production Systems, books in Springer Series in Operations Research/Computer Science Interfaces. The challenge with this work is to gather the main contribution in three fields, optimization technique for production decision, general development for optimization and computing method and wider spread applications. The number of researches dealing with decision maker tool and optimization method grows very quickly these last years and in a large number of fields. We may be able to read nice and worthy works from research developed in chemical, mechanical, computing, automotive and many other fields.

How do you manage deployment of IoT applications? Which operating systems are running OpenStack deployments? Is RFP requirement that login access to the platform is differentiated by tenant or should all logs be created differently per tenant? What are your edge computing characteristics, use cases, and scenarios? What container and PaaS tools are used to manage OpenStack applications? This breakthrough OpenStack self-assessment will make you the accepted OpenStack domain veteran by revealing just what you need to know to be fluent and ready for any OpenStack challenge. How do I reduce the effort in the OpenStack work to be done to get problems solved? How can I ensure that plans of action include every OpenStack task and that every OpenStack outcome is in place? How will I save time investigating strategic and tactical options and ensuring OpenStack costs are low? How can I deliver tailored OpenStack advice instantly with structured going-forward plans? There's no better guide through these mind-expanding questions than acclaimed best-selling author Gerard Blokdyk. Blokdyk ensures all OpenStack essentials are covered, from every angle: the OpenStack self-assessment shows succinctly and clearly that what needs to be clarified to organize the required activities and processes so that OpenStack outcomes are achieved. Contains extensive criteria grounded in past and current successful projects and activities by experienced OpenStack practitioners. Their mastery, combined with the easy elegance of the self-assessment, provides its superior value to you in knowing how to ensure the outcome of any efforts in OpenStack are maximized with professional results. Your purchase includes access details to the OpenStack self-assessment dashboard download which gives you your dynamically prioritized projects-ready tool and shows you exactly what to

do next. Your exclusive instant access details can be found in your book. You will receive the following contents with New and Updated specific criteria: - The latest quick edition of the book in PDF - The latest complete edition of the book in PDF, which criteria correspond to the criteria in... - The Self-Assessment Excel Dashboard - Example pre-filled Self-Assessment Excel Dashboard to get familiar with results generation - In-depth and specific OpenStack Checklists - Project management checklists and templates to assist with implementation INCLUDES LIFETIME SELF ASSESSMENT UPDATES Every self assessment comes with Lifetime Updates and Lifetime Free Updated Books. Lifetime Updates is an industry-first feature which allows you to receive verified self assessment updates, ensuring you always have the most accurate information at your fingertips.

Openstack in ActionManning Publications Company

Summary Meteor in Action teaches you full-stack web development using the Meteor platform. It starts with an overview of a Meteor application, revealing the unique nature of Meteor's end-to-end application model. Then you'll dive into the Blaze templating engine, discover Meteor's reactive data sources model, learn simple and advanced routing techniques, and practice managing users, permissions, and roles. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book You might call Meteor a reactive, isomorphic, full-stack web development framework. Or, like most developers who have tried it, you might just call it awesome. Meteor is a JavaScript-based framework for both client and server web and mobile applications. Meteor applications react to changes in data instantly, so you get impossibly responsive user experiences, and the consistent build process, unified front- and back-end package system, and one-command deploys save you time at every step from design to release. Meteor in Action teaches you full-stack web development with Meteor. It starts by revealing the unique nature of Meteor's end-to-end application model. Through real-world scenarios, you'll dive into the Blaze templating engine, discover Meteor's reactive data sources model, learn routing techniques, and practice managing users, permissions, and roles. Finally, you'll learn how to deploy Meteor on your server and scale efficiently. What's Inside Building your first real-time application Using MongoDB and other reactive data sources Creating applications with Iron Router Deploying and scaling your applications About the Reader Readers need to know the basics of JavaScript and understand general web application design. About the Authors Stephan Hochhaus and Manuel Schoebel are veteran web developers who have worked with Meteor since its infancy. Table of Contents PART 1 LOOK—A SHOOTING STAR! A better way to build apps My fridge! A reactive gamePART 2 3, 2, 1—IMPACT! Working with templates Working with data Fully reactive editing Users, authentications, and permissions Exchanging data Routing using Iron.Router The package system Advanced server methods PART 3 LEAVING THE CRATER Building and debugging Going into production

Summary Mesos in Action introduces readers to the Apache Mesos cluster manager and the concept of application-centric infrastructure. Filled with helpful figures and hands-on instructions, this book guides you from your first steps creating a highly-available Mesos cluster through deploying applications in production and writing native Mesos frameworks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Modern datacenters are complex environments, and when you throw Docker and other container-based systems into the mix, there's a great need to simplify. Mesos is an open source cluster management platform that transforms the whole datacenter into a single pool of compute, memory, and storage resources that you can allocate, automate, and scale as if you're working with a single supercomputer. About the Book Mesos in Action introduces readers to the Apache Mesos cluster manager and the concept of application-centric infrastructure. Filled with helpful figures and hands-on instructions, this book guides you from your first steps creating a highly-available Mesos cluster through deploying applications in production and writing native Mesos frameworks. You'll learn how to scale to thousands of nodes, while providing resource isolation between processes using Linux and Docker containers. You'll also learn practical techniques for deploying applications using popular key frameworks. What's Inside Spinning up your first Mesos cluster Scheduling, resource administration, and logging Deploying containerized applications with Marathon, Chronos, and Aurora Writing Mesos frameworks using Python About the Reader Readers need to be familiar with the core ideas of datacenter administration and need a basic knowledge of Python or a similar programming language. About the Author Roger Ignazio is an experienced systems engineer with a focus on distributed, fault-tolerant, and scalable infrastructure. He is currently a technical lead at Mesosphere. Table of Contents PART 1 HELLO, MESOS Introducing Mesos Managing datacenter resources with Mesos PART 2 CORE MESOS Setting up Mesos Mesos fundamentals Logging and debugging Mesos in production PART 3 RUNNING ON MESOS Deploying applications with MarathoN Managing scheduled tasks with Chronos Deploying applications and managing scheduled tasks with Aurora Developing a framework

Get up and running with OpenStack Swift, the free, open source solution for deploying high-performance object storage clusters at scale. In this practical guide, Joe Arnold, co-founder and CEO of SwiftStack, brings you up-to-speed on the basic concepts of object storage and walks you through what you need to know to plan, build, operate, and measure the performance of your own Swift storage system. Object storage is essential today with the growth of web, mobile, and software-as-a-service (SaaS) applications. This book helps you through the process, with separate sections on application development, installation, administration, and troubleshooting. Learn Swift's concepts for organizing, distributing, and serving data Explore basic and advanced features of the Swift RESTful API Delve into Swift's many client libraries, including useful Python features Write middleware to customize and simplify your storage system Understand requirements for planning a Swift deployment—including your specific use case Learn options for coaxing the best performance from your cluster Get best practices for daily operations, such as monitoring and planning capacity additions Pick up techniques for testing and benchmarking your Swift cluster OpenStack Trove is your step-by-step guide to set up and run a secure and scalable cloud Database as a Service (DBaaS) solution. The book shows you how to set up and configure the Trove DBaaS framework, use prepackaged or custom database implementations, and provision and operate a variety of databases—including MySQL, PostgreSQL, MongoDB, Cassandra, and Redis—in development and production environments. Authors Amrith Kumar and Douglas Shelley, both active technical contributors to the Trove project, describe common deployment scenarios such as single-node database instances and walk you through the setup, configuration, and ongoing management of complex database topics like replication, clustering, and high availability. The book provides detailed descriptions of how Trove works and gives you an in-depth understanding of its architecture. It also shows you how to avoid common errors and debug and troubleshoot Trove installations, and perform common tasks such as:

OpenStack is a system that controls large pools of computing, storage, and networking resources, allowing its users to provision resources through a user-friendly interface. OpenStack helps developers with features such as rolling upgrades, federated identity,

and software reliability. You will begin with basic security policies, such as MAC, MLS, and MCS, and explore the structure of OpenStack and virtual networks with Neutron. Next, you will configure secure communications on the OpenStack API with HTTP connections. You will also learn how to set OpenStack Keystone and OpenStack Horizon and gain a deeper understanding of the similarities/differences between OpenStack Cinder and OpenStack Swift. By the end of this book, you will be able to tweak your hypervisor to make it safer and a smart choice based on your needs.

Discover the basics of virtual networking in OpenStack to implement various cloud network architectures Key Features Learn the difference between Open vSwitch and Linux bridge switching technologies Connect virtual machine instances to virtual networks, subnets, and ports Implement virtual load balancers, firewalls, and routers in your network Book Description OpenStack Networking is a pluggable, scalable, and API-driven system to manage physical and virtual networking resources in an OpenStack-based cloud. Like other core OpenStack components, OpenStack Networking can be used by administrators and users to increase the value and maximize the use of existing datacenter resources. This third edition of Learning OpenStack Networking walks you through the installation of OpenStack and provides you with a foundation that can be used to build a scalable and production-ready OpenStack cloud. In the initial chapters, you will review the physical network requirements and architectures necessary for an OpenStack environment that provide core cloud functionality. Then, you'll move through the installation of the new release of OpenStack using packages from the Ubuntu repository. An overview of Neutron networking foundational concepts, including networks, subnets, and ports will segue into advanced topics such as security groups, distributed virtual routers, virtual load balancers, and VLAN tagging within instances. By the end of this book, you will have built a network infrastructure for your cloud using OpenStack Neutron. What you will learn Get familiar with Neutron constructs, including agents and plugins Build foundational Neutron resources to provide connectivity to instances Work with legacy Neutron routers and troubleshoot traffic through them Explore high-availability routing capabilities utilizing Virtual Router Redundancy Protocol (VRRP) Create and manage load balancers and associated components Manage security groups as a method of securing traffic to and from instances Who this book is for If you are an OpenStack-based cloud operator and administrator who is new to Neutron networking and wants to build your very own OpenStack cloud, then this book is for you. Prior networking experience and a physical server and network infrastructure is recommended to follow along with concepts demonstrated in the book.

Over 90 practical and highly applicable recipes to successfully deploy various OpenStack configurations in production About This Book Get a deep understanding of OpenStack's internal structure and services Learn real-world examples on how to build and configure various production grade use cases for each of OpenStack's services Use a step-by-step approach to install and configure OpenStack's services to provide Compute, Storage, and Networking as a services for cloud workloads Who This Book Is For If you have a basic understanding of Linux and Cloud computing and want to learn about configurations that OpenStack supports, this is the book for you. Knowledge of virtualization and managing Linux environments is expected. Prior knowledge or experience of OpenStack is not required, although beneficial. What You Will Learn Plan an installation of OpenStack with a basic configuration Deploy OpenStack in a highly available configuration Configure Keystone Identity services with multiple types of identity backends Configure Glance Image Store with File, NFS, Swift, or Ceph image backends and use local image caching Design Cinder to use a single storage provider such as LVM, Ceph, and NFS backends, or to use multiple storage backends simultaneously Manage and configure the OpenStack networking backend Configure OpenStack's compute hypervisor and the instance scheduling mechanism Build and customize the OpenStack dashboard In Detail OpenStack is the most popular open source cloud platform used by organizations building internal private clouds and by public cloud providers. OpenStack is designed in a fully distributed architecture to provide Infrastructure as a Service, allowing us to maintain a massively scalable cloud infrastructure. OpenStack is developed by a vibrant community of open source developers who come from the largest software companies in the world. The book provides a comprehensive and practical guide to the multiple uses cases and configurations that OpenStack supports. This book simplifies the learning process by guiding you through how to install OpenStack in a single controller configuration. The book goes deeper into deploying OpenStack in a highly available configuration. You'll then configure Keystone Identity Services using LDAP, Active Directory, or the MySQL identity provider and configure a caching layer and SSL. After that, you will configure storage back-end providers for Glance and Cinder, which will include Ceph, NFS, Swift, and local storage. Then you will configure the Neutron networking service with provider network VLANs, and tenant network VXLAN and GRE. Also, you will configure Nova's Hypervisor with KVM, and QEMU emulation, and you will configure Nova's scheduler filters and weights. Finally, you will configure Horizon to use Apache HTTPD and SSL, and you will customize the dashboard's appearance. Style and approach This book consists of clear, concise instructions coupled with practical and applicable recipes that will enable you to use and implement the latest features of OpenStack.

Summary OpenStack in Action offers the real world use cases and step-by-step instructions you can take to develop your own cloud platform from from inception to deployment. This book guides you through the design of both the physical hardware cluster and the infrastructure services you'll need to create a custom cloud platform. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology OpenStack is an open source framework that lets you create a private or public cloud platform on your own physical servers. You build custom infrastructure, platform, and software services without the expense and vendor lock-in associated with proprietary cloud platforms like Amazon Web Services and Microsoft Azure. With an OpenStack private cloud, you can get increased security, more control, improved reliability, and lower costs. About the Book OpenStack in Action offers real-world use cases and step-by-step instructions on how to develop your own cloud platform. This book guides you through the design of both the physical hardware cluster and the infrastructure services you'll need. You'll learn how to select and set up virtual and physical servers, how to implement software-defined networking, and technical details of designing, deploying, and operating an OpenStack cloud in your enterprise. You'll also discover how to best tailor your OpenStack deployment for your environment. Finally, you'll learn how your cloud can offer user-facing software and infrastructure services.

What's Inside Develop and deploy an enterprise private cloud Private cloud technologies from an IT perspective Organizational impact of self-service cloud computing About the Reader No prior knowledge of OpenStack or cloud development is assumed. About the Author Cody Bumgardner is the Chief Technology Architect at a large university where he is responsible for the architecture, deployment, and long-term strategy of OpenStack private clouds and other cloud computing initiatives. Table of Contents PART 1 GETTING STARTED Introducing OpenStack Taking an OpenStack test-drive Learning basic OpenStack operations Understanding private cloud building blocks PART 2 WALKING THROUGH A MANUAL DEPLOYMENT Walking through a Controller deployment Walking through a Networking deployment Walking through a Block Storage deployment Walking through a Compute deployment PART 3 BUILDING A PRODUCTION ENVIRONMENT Architecting your OpenStack Deploying Ceph Automated HA OpenStack deployment with Fuel Cloud orchestration using OpenStack

A Cookbook full of practical and applicable recipes that will enable you to use the full capabilities of OpenStack like never before. This book is aimed at system administrators and technical architects moving from a virtualized environment to cloud environments with familiarity of cloud computing platforms. Knowledge of virtualization and managing linux environments is expected.

Summary Cloud Native Patterns is your guide to developing strong applications that thrive in the dynamic, distributed, virtual world of the cloud. This book presents a mental model for cloud-native applications, along with the patterns, practices, and tooling that set them apart. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Cloud platforms promise the holy grail: near-zero downtime, infinite scalability, short feedback cycles, fault-tolerance, and cost control. But how do you get there? By applying cloud-native designs, developers can build resilient, easily adaptable, web-scale distributed applications that handle massive user traffic and data loads. Learn these fundamental patterns and practices, and you'll be ready to thrive in the dynamic, distributed, virtual world of the cloud. About the Book With 25 years of experience under her belt, Cornelia Davis teaches you the practices and patterns that set cloud-native applications apart. With realistic examples and expert advice for working with apps, data, services, routing, and more, she shows you how to design and build software that functions beautifully on modern cloud platforms. As you read, you will start to appreciate that cloud-native computing is more about the how and why rather than the where. What's inside The lifecycle of cloud-native apps Cloud-scale configuration management Zero downtime upgrades, versioned services, and parallel deploys Service discovery and dynamic routing Managing interactions between services, including retries and circuit breakers About the Reader Requires basic software design skills and an ability to read Java or a similar language. About the Author Cornelia Davis is Vice President of Technology at Pivotal Software. A teacher at heart, she's spent the last 25 years making good software and great software developers. Table of Contents PART 1 - THE CLOUD-NATIVE CONTEXT You keep using that word: Defining "cloud-native" Running cloud-native applications in production The platform for cloud-native software PART 2 - CLOUD-NATIVE PATTERNS Event-driven microservices: It's not just request/response App redundancy: Scale-out and statelessness Application configuration: Not just environment variables The application lifecycle: Accounting for constant change Accessing apps: Services, routing, and service discovery Interaction redundancy: Retries and other control loops Fronting services: Circuit breakers and API gateways Troubleshooting: Finding the needle in the haystack Cloud-native data: Breaking the data monolith

Design and implement successful private clouds with OpenStack About This Book Explore the various design choices available for cloud architects within an OpenStack deployment Craft an OpenStack architecture and deployment pipeline to meet the unique needs of your organization Create a product roadmap for Infrastructure as a Service in your organization using this hands-on guide Who This Book Is For This book is written especially for those who will design OpenStack clouds and lead their implementation. These people are typically cloud architects, but may also be in product management, systems engineering, or enterprise architecture. What You Will Learn Familiarize yourself with the components of OpenStack Build an increasingly complex OpenStack lab deployment Write compelling documentation for the architecture teams within your organization Apply Agile configuration management techniques to deploy OpenStack Integrate OpenStack with your organization's identity management, provisioning, and billing systems Configure a robust virtual environment for users to interact with Use enterprise security guidelines for your OpenStack deployment Create a product roadmap that delivers functionality quickly to the users of your platform In Detail Over the last five years, hundreds of organizations have successfully implemented Infrastructure as a Service (IaaS) platforms based on OpenStack. The huge amount of investment from these organizations, industry giants such as IBM and HP, as well as open source leaders such as Red Hat have led analysts to label OpenStack as the most important open source technology since the Linux operating system. Because of its ambitious scope, OpenStack is a complex and fast-evolving open source project that requires a diverse skill-set to design and implement it. This guide leads you through each of the major decision points that you'll face while architecting an OpenStack private cloud for your organization. At each point, we offer you advice based on the experience we've gained from designing and leading successful OpenStack projects in a wide range of industries. Each chapter also includes lab material that gives you a chance to install and configure the technologies used to build production-quality OpenStack clouds. Most importantly, we focus on ensuring that your OpenStack project meets the needs of your organization, which will guarantee a successful rollout. Style and approach This is practical, hands-on guide to implementing OpenStack clouds, where each topic is illustrated with real-world examples and then the technical points are proven in the lab.

Design, build, and automate 10 real-world OpenStack administrative tasks with Ansible About This Book Automate real-world OpenStack cloud operator administrative tasks Construct a collection of automation code to save time on managing your OpenStack cloud Use this step-by-step tutorial to automate such tasks with Ansible Who This Book Is For

If you are an OpenStack-based cloud operator and/or infrastructure administrator and are interested in automating administrative functions, then this book is exactly what you are looking for. Having a functioning OpenStack environment is helpful, but most certainly not necessary. What You Will Learn Efficiently execute OpenStack administrative tasks Familiarize yourself with how Ansible works and assess the defined best practices Create Ansible playbooks and roles Automate tasks to customize your OpenStack cloud Review OpenStack automation considerations when automating administrative tasks Examine and automate advanced OpenStack tasks and designated use cases Get a high-level overview of OpenStack and the current production-ready projects Deep dive into OpenStack CLI tools and find out how to use them In Detail Most organizations are seeking methods to improve business agility because they have realized just having a cloud is not enough. Being able to improve application deployments, reduce infrastructure downtime, and eliminate daily manual tasks can only be accomplished through some sort of automation. Packed with real-world OpenStack administrative tasks, this book will walk you through working examples and explain how these tasks can be automated using one of the most popular open source automation tools—Ansible. We will start with a brief overview of OpenStack and Ansible and highlight some best practices. Each chapter will provide an introduction to handling various Cloud Operator administration tasks such as creating multiple users/tenants, setting up Multi-Tenant Isolation, customizing your clouds quotas, taking instance snapshots, evacuating compute hosts for maintenance, and running cloud health checks, and a step-by-step tutorial on how to automate these tasks with Ansible. Style and approach This easy-to-follow reference guide is packed with examples of real-world OpenStack administration tasks; each task is explained in detail and then subsequently turned into automation code.

If you need to get started with OpenStack or want to learn more, then this book is your perfect companion. If you're comfortable with the Linux command line, you'll gain confidence in using OpenStack.

Your complete guide to designing, deploying, and managing OpenStack-based clouds in mid-to-large IT infrastructures About This Book* Design and deploy an OpenStack-based cloud in your mid-to-large IT infrastructure using automation tools and best practices* Keep yourself up-to-date with valuable insights into OpenStack components and new services in the latest OpenStack release* Discover how the new features in the latest OpenStack release can help your enterprise and infrastructure Who This Book Is For This book is for system administrators, cloud engineers, and system architects who would like to deploy an OpenStack-based cloud in a mid-to-large IT infrastructure. This book requires a moderate level of system administration and familiarity with cloud concepts. What You Will Learn* Explore the main architecture design of OpenStack components and core-by-core services, and how they work together* Design different high availability scenarios and plan for a no-single-point-of-failure environment* Set up a multinode environment in production using orchestration tools* Boost OpenStack's performance with advanced configuration* Delve into various hypervisors and container technology supported by OpenStack* Get familiar with deployment methods and discover use cases in a real production environment* Adopt the DevOps style of automation while deploying and operating in an OpenStack environment* Monitor the cloud infrastructure and make decisions on maintenance and performance improvement In Detail In this second edition, you will get to grips with the latest features of OpenStack. Starting with an overview of the OpenStack architecture, you'll see how to adopt the DevOps style of automation while deploying and operating in an OpenStack environment. We'll show you how to create your own OpenStack private cloud. Then you'll learn about various hypervisors and container technology supported by OpenStack. You'll get an understanding about the segregation of compute nodes based on reliability and availability needs. We'll cover various storage types in OpenStack and advanced networking aspects such as SDN and NFV. Next, you'll understand the OpenStack infrastructure from a cloud user point of view. Moving on, you'll develop troubleshooting skills, and get a comprehensive understanding of services such as high availability and failover in OpenStack. Finally, you will gain experience of running a centralized logging server and monitoring OpenStack services. The book will show you how to carry out performance tuning based on OpenStack service logs. You will be able to master OpenStack benchmarking and performance tuning. By the end of the book, you'll be ready to take steps to deploy and manage an OpenStack cloud with the latest open source technologies.

Summary OpenStack in Action offers the real world use cases and step-by-step instructions you can take to develop your own cloud platform from from inception to deployment. This book guides you through the design of both the physical hardware cluster and the infrastructure services you'll need to create a custom cloud platform. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology OpenStack is an open source framework that lets you create a private or public cloud platform on your own physical servers. You build custom infrastructure, platform, and software services without the expense and vendor lock-in associated with proprietary cloud platforms like Amazon Web Services and Microsoft Azure. With an OpenStack private cloud, you can get increased security, more control, improved reliability, and lower costs. About the Book "OpenStack in Action" offers real-world use cases and step-by-step instructions on how to develop your own cloud platform. This book guides you through the design of both the physical hardware cluster and the infrastructure services you'll need. You'll learn how to select and set up virtual and physical servers, how to implement software-defined networking, and technical details of designing, deploying, and operating an OpenStack cloud in your enterprise. You'll also discover how to best tailor your OpenStack deployment for your environment. Finally, you'll learn how your cloud can offer user-facing software and infrastructure services.

What's Inside Develop and deploy an enterprise private cloud Private cloud technologies from an IT perspective Organizational impact of self-service cloud computing About the Reader No prior knowledge of OpenStack or cloud development is assumed. About the Author Cody Bumgardner is the Chief Technology Architect at a large university where he is responsible for the architecture, deployment, and long-term strategy of OpenStack private clouds and other cloud computing initiatives. Table of Contents PART 1 GETTING STARTED Introducing OpenStack Taking an

OpenStack test-drive Learning basic OpenStack operations Understanding private cloud building blocks PART 2 WALKING THROUGH A MANUAL DEPLOYMENT Walking through a Controller deployment Walking through a Networking deployment Walking through a Block Storage deployment Walking through a Compute deployment PART 3 BUILDING A PRODUCTION ENVIRONMENT Architecting your OpenStack Deploying Ceph Automated HA OpenStack deployment with Fuel Cloud orchestration using OpenStack

This IBM® Redbooks® publication introduces the IBM Software Defined Environment (SDE) solution, which helps to optimize the entire computing infrastructure--compute, storage, and network resources--so that it can adapt to the type of work required. In today's environment, resources are assigned manually to workloads, but that happens automatically in a SDE. In an SDE, workloads are dynamically assigned to IT resources based on application characteristics, best-available resources, and service level policies so that they deliver continuous, dynamic optimization and reconfiguration to address infrastructure issues. Underlying all of this are policy-based compliance checks and updates in a centrally managed environment. Readers get a broad introduction to the new architecture. Think integration, automation, and optimization. Those are enablers of cloud delivery and analytics. SDE can accelerate business success by matching workloads and resources so that you have a responsive, adaptive environment. With the IBM Software Defined Environment, infrastructure is fully programmable to rapidly deploy workloads on optimal resources and to instantly respond to changing business demands. This information is intended for IBM sales representatives, IBM software architects, IBM Systems Technology Group brand specialists, distributors, resellers, and anyone who is developing or implementing SDE.

Wield the power of OpenStack Neutron networking to bring network infrastructure and capabilities to your cloud About This Book This completely up-to-date edition will show you how to deploy a cloud on OpenStack using community-driven processes. It includes rich examples that will help you understand complex networking topics with ease Understand every aspect of designing, creating, customizing, and maintaining the core network foundation of an OpenStack cloud using OpenStack Neutron all in one book Written by best-selling author James Denton, who has more than 15 years of experience in system administration and networking. James has experience of deploying, operating, and maintaining OpenStack clouds and has worked with top enterprises and organizations Who This Book Is For If you are an OpenStack-based cloud operator and administrator who is new to Neutron networking and wants to build your very own OpenStack cloud, then this book is for you. Prior networking experience and a physical server and network infrastructure is recommended to follow along with concepts demonstrated in the book. What You Will Learn Architect and install the latest release of OpenStack on Ubuntu Linux 14.04 LTS Review the components of OpenStack networking, including plugins, agents, and services, and learn how they work together to coordinate network operations Build a virtual switching infrastructure using reference architectures based on ML2 + Open vSwitch or ML2 + LinuxBridge Create networks, subnets, and routers that connect virtual machine instances to the network Deploy highly available routers using DVR or VRRP-based methods Scale your application with haproxy and Load Balancing as-a-Service Implement port and router-level security using Security Groups and Firewall as-a-Service Provide connectivity to tenant networks with Virtual Private Networking as-a-Service (VPNaaS) Find out how to manage OpenStack networking resources using CLI and GUI-driven methods In Detail OpenStack Neutron is an OpenStack component that provides networking as a service for other OpenStack services to architect networks and create virtual machines through its API. This API lets you define network connectivity in order to leverage network capabilities to cloud deployments. Through this practical book, you will build a strong foundational knowledge of Neutron, and will architect and build an OpenStack cloud using advanced networking features. We start with an introduction to OpenStack Neutron and its various components, including virtual switching, routing, FWaaS, VPNaaS, and LBaaS. You'll also get hands-on by installing OpenStack and Neutron and its components, and use agents and plugins to orchestrate network connectivity and build a virtual switching infrastructure. Moving on, you'll get to grips with the HA routing capabilities utilizing VRRP and distributed virtual routers in Neutron. You'll also discover load balancing fundamentals, including the difference between nodes, pools, pool members, and virtual IPs. You'll discover the purpose of security groups and learn how to apply the security concept to your cloud/tenant/instance. Finally, you'll configure virtual private networks that will allow you to avoid the use of SNAT and floating IPs when connecting to remote networks. Style and approach This easy-to-follow guide on networking in OpenStack follows a step-by-step process to installing OpenStack and configuring the base networking components. Each major networking component has a dedicated chapter that will build on your experience gained from prior chapters. Set up and maintain your own cloud-based Infrastructure as a Service (IaaS) using OpenStack About This Book • Build and manage a cloud environment using just four virtual machines • Get to grips with mandatory as well as optional OpenStack components and know how they work together • Leverage your cloud environment to provide Infrastructure as a Service (IaaS) with this practical, step-by-step guide Who This Book Is For This book is targeted at all aspiring administrators, architects, or students who want to build cloud environments using Openstack. Knowledge of IaaS or cloud computing is recommended. What You Will Learn • Get an introduction to OpenStack and its components • Authenticate and authorize the cloud environment using Keystone • Store and retrieve data and images using storage components such as Cinder, Swift, and Glance • Use Nova to build a Cloud Computing fabric controller • Abstract technology-agnostic networks using the Neutron network component • Gain an understanding of optional components such as Ceilometer, Trove, Ironic, Sahara, Barbican, Zaqar, Designate, Manila, and many more • See how all of the OpenStack components collaborate to provide IaaS to users • Create a production-grade OpenStack and automate your OpenStack Cloud In Detail OpenStack is a free and open source cloud computing platform that is rapidly gaining popularity in Enterprise data centres. It is a scalable operating system and is used to build private and public clouds. It is imperative for all the aspiring cloud administrators to possess OpenStack skills if they want to succeed in the cloud-led IT

infrastructure space. This book will help you gain a clearer understanding of OpenStack's components and their interaction with each other to build a cloud environment. You will learn to deploy a self-service based cloud using just four virtual machines and standard networking. You begin with an introduction on the basics of cloud computing. This is followed by a brief look into the need for authentication and authorization, the different aspects of dashboards, cloud computing fabric controllers, along with "Networking as a Service" and "Software Defined Networking." Then, you will focus on installing, configuring, and troubleshooting different architectures such as Keystone, Horizon, Nova, Neutron, Cinder, Swift, and Glance. Furthermore, you will see how all of the OpenStack components come together in providing IaaS to users. Finally, you will take your OpenStack cloud to the next level by integrating it with other IT ecosystem elements before automation. By the end of this book, you will be proficient with the fundamentals and application of OpenStack.

Style and approach This is a practical step-by-step guide comprising of installation prerequisites and basic troubleshooting instructions to help you build an error-free OpenStack cloud easily.

Summary CoreOS in Action is a clear tutorial for deploying container-based systems on CoreOS Container Linux. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Traditional Linux server distributions include every component required for anything you might be hosting, most of which you don't need if you've containerized your apps and services. CoreOS Container Linux is a bare-bones distro with only the essential bits needed to run containers like Docker. Container Linux is a fast, secure base layer for any container-centric distributed application, including microservices. And say goodbye to patch scheduling; when Container Linux needs an update, you just hot-swap the whole OS. About the Book CoreOS in Action is a clear tutorial for deploying container-based systems on CoreOS Container Linux. Inside, you'll follow along with examples that teach you to set up CoreOS on both private and cloud systems, and to practice common sense monitoring and upgrade techniques with real code. You'll also explore important container-aware application designs, including microservices, web, and Big Data examples with real-world use cases to put your learning into perspective. Summary Handling scaling and failures gracefully Container-driven application designs Cloud, on-premises, and hybrid deployment Smart logging and backup practices About the Reader Written for readers familiar with Linux and the basics of Docker. About the Author Matt Bailey is currently a technical lead at ZeniMax. He has worked in higher education and with scientific computing, medical, and networking technology companies, as well as a few startups. You can find him online via <https://mdb.io>.

Table of Contents

PART 1 - GETTING TO KNOW COREOS Introduction to the CoreOS family Getting started on your workstation Expecting failure: fault tolerance in CoreOS

PART 2 - APPLICATION ARCHITECTURE CoreOS in production Application architecture and workflow Web stack application example Big Data stack

PART 3 - COREOS IN PRODUCTION CoreOS on AWS Bringing it together: deployment System administration

Develop and manage applications on the AWS and OpenStack platforms with this comprehensive learning guide. Key Features A step-by-step guide to help you develop applications on the hybrid cloud platform. Acquire an in-depth understanding of the OpenStack and AWS cloud platforms. Extensive source code examples for OpenStack and AWS applications. Easily troubleshoot OpenStack and AWS issues. Understand the best practices and security measures for the hybrid cloud platform. Book Description This book introduces you to the hybrid cloud platform, and focuses on the AWS public cloud and OpenStack private cloud platforms. It provides a deep dive into the AWS and OpenStack cloud platform services that are essential for developing hybrid cloud applications. You will learn to develop applications on AWS and OpenStack platforms with ease by leveraging various cloud services and taking advantage of PaaS. The book provides you with the ability to leverage the flexibility of choosing a cloud platform for migrating your existing resources to the cloud, as well as developing hybrid cloud applications that can migrate virtual machine instances from AWS to OpenStack and vice versa. You will also be able to build and test cloud applications without worrying about the system that your development environment supports. The book also provides an in-depth understanding of the best practices that are followed across the industry for developing cloud applications, as well as for adapting the hybrid cloud platform. Lastly, it also sheds light on various troubleshooting techniques for OpenStack and AWS cloud platform services that are consumed by hybrid cloud applications. By the end of this book, you will have a deep understanding of the hybrid cloud platform and will be able to develop robust, efficient, modular, scalable, and flexible cloud applications. What you will learn Understand the hybrid cloud platform Explore the AWS and OpenStack cloud platforms in depth Develop AWS applications with source code examples Develop OpenStack applications with source code examples Troubleshoot OpenStack and AWS Learn hybrid cloud best practices Understand security measures on the hybrid cloud Who this book is for If you are an IT professional, developer, or a DevOps engineer looking to develop and manage your applications on the hybrid cloud platform, then this book is for you. Some prior knowledge of the public and private cloud will enhance your skills. Developers looking to build applications using AWS or OpenStack services will also benefit from this book.

Summary Amazon Web Services in Action, Second Edition is a comprehensive introduction to computing, storing, and networking in the AWS cloud. You'll find clear, relevant coverage of all the essential AWS services you to know, emphasizing best practices for security, high availability and scalability. Foreword by Ben Whaley, AWS community hero and author. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The largest and most mature of the cloud platforms, AWS offers over 100 prebuilt services, practically limitless compute resources, bottomless secure storage, as well as top-notch automation capabilities. This book shows you how to develop, host, and manage applications on AWS. About the Book Amazon Web Services in Action, Second Edition is a comprehensive introduction to deploying web applications in the AWS cloud. You'll find clear, relevant coverage of all essential AWS services, with a focus on automation, security, high availability, and scalability. This thoroughly revised edition covers the latest additions to AWS, including serverless infrastructure with AWS Lambda, sharing data with EFS, and in-memory storage with ElastiCache. What's inside Completely revised

bestseller Secure and scale distributed applications Deploy applications on AWS Design for failure to achieve high availability Automate your infrastructure About the Reader Written for mid-level developers and DevOps engineers. About the Author Andreas Wittig and Michael Wittig are software engineers and DevOps consultants focused on AWS. Together, they migrated the first bank in Germany to AWS in 2013. Table of Contents PART 1 - GETTING STARTED What is Amazon Web Services? A simple example: WordPress in five minutes PART 2 - BUILDING VIRTUAL INFRASTRUCTURE CONSISTING OF COMPUTERS AND NETWORKING Using virtual machines: EC2 Programming your infrastructure: The command-line, SDKs, and CloudFormation Automating deployment: CloudFormation, Elastic Beanstalk, and OpsWorks Securing your system: IAM, security groups, and VPC Automating operational tasks with Lambda PART 3 - STORING DATA IN THE CLOUD Storing your objects: S3 and Glacier Storing data on hard drives: EBS and instance store Sharing data volumes between machines: EFS Using a relational database service: RDS Caching data in memory: Amazon ElastiCache Programming for the NoSQL database service: DynamoDB PART 4 - ARCHITECTING ON AWS Achieving high availability: availability zones, auto-scaling, and CloudWatch Decoupling your infrastructure: Elastic Load Balancing and Simple Queue Service Designing for fault tolerance Scaling up and down: auto-scaling and CloudWatch

[Copyright: 0a56c3fa58262a7f6276eb8182f54a23](#)