

Thank you for reading Openstack Api Documentation. As you may know, people have look numerous times for their favorite novels like this Openstack Api Documentation, but end up in infectious downloads.

Rather than reading a good book with a cup of tea in the afternoon, instead they are facing with some infectious bugs inside their laptop.

Openstack Api Documentation is available in our book collection an online access to it is set as public so you can get it instantly.

Our digital library saves in multiple countries, allowing you to get the most less latency time to download any of our books like this one.

Merely said, the Openstack Api Documentation is universally compatible with any devices to read



[OpenStack for Architects](#) Springer Nature Become an expert in implementing advanced, network-related tasks with Python. About This Book Build the skills to perform all networking tasks using Python with ease Use Python for network device automation, DevOps, and software-defined networking Get practical guidance to networking with Python Who This Book Is For If you are a network engineer or a programmer who wants to use Python for networking, then this book is for you. A basic familiarity with networking-related concepts such as TCP/IP and a familiarity with Python programming will be useful. What You Will Learn Review all the fundamentals of Python and the TCP/IP suite Use Python to execute commands when the device does not support the API or programmatic interaction with the device Implement automation techniques by integrating Python with Cisco, Juniper, and Arista eAPI Integrate Ansible using Python to control Cisco, Juniper, and Arista networks Achieve network security with Python Build Flask-based web-service APIs with Python Construct a Python-based migration plan from a legacy to scalable SDN-based network. In Detail This book begins with a review of the TCP/ IP protocol suite and a refresher of the core elements of the Python language. Next, you will start using Python and supported libraries to automate network tasks from the current major network vendors. We will look at automating traditional network devices based on the command-line interface, as well as newer devices with API support, with hands-on labs. We will then learn the concepts and practical use cases of the Ansible framework in order to achieve your network goals. We will then move on to using Python for DevOps, starting with using open source tools to test, secure, and analyze your network. Then, we will focus on network monitoring and visualization. We will learn how to retrieve network information using a polling mechanism, ?ow-based monitoring, and visualizing the data programmatically. Next, we will learn how to use the Python framework to build your own customized network web services. In the last module, you will use Python for SDN, where you will use a Python-based controller with OpenFlow in a hands-on lab to learn its concepts and applications. We will compare and contrast OpenFlow, OpenStack, OpenDaylight, and NFV. Finally, you will use everything you've learned in the book to construct a migration plan to go from a legacy to a scalable SDN-based network. Style and approach An easy-to-follow guide packed with hands-on examples of using Python for network device automation, DevOps, and SDN. IBM Cloud Manager with OpenStack on z Systems V4.2 Packt Publishing Ltd

Throughout human history, technological advancements have been made for the ease of human labor. With our most recent advancements, it has been the work of scholars to discover ways for machines to take over a large part of this labor and reduce human intervention. These advancements may become essential processes

to nearly every industry. It is essential to be knowledgeable about automation so that it may be applied. Research Anthology on Cross-Disciplinary Designs and Applications of Automation is a comprehensive resource on the emerging designs and application of automation. This collection features a number of authors spanning multiple disciplines such as home automation, healthcare automation, government automation, and more. Covering topics such as human-machine interaction, trust calibration, and sensors, this research anthology is an excellent resource for technologists, IT specialists, computer engineers, systems and software engineers, manufacturers, engineers, government officials, professors, students, healthcare administration, managers, CEOs, researchers, and academicians.

[IBM PowerKVM: Configuration and Use](#) Springer

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

[Mastering Python Networking](#) "O'Reilly Media, Inc."

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.

[Manage Your Clouds with IBM Cloud Manager with OpenStack for z Systems, V4.2](#) Apress

Looking for a way to invigorate your technical writing team and grow that expertise to include developers, designers, and writers of all backgrounds? When you treat docs like code, you multiply everyone's efforts and streamline processes through collaboration, automation, and innovation. Second edition now available with updates and more information about version control for documents and continuous publishing.

[OpenStack: Building a Cloud Environment](#) Packt Publishing Ltd

This book constitutes the proceedings of the 14th International Conference on Service-Oriented Computing, ICSOC 2016, held in Banff, AB, Canada, in October 2016. The 30 full papers presented together with 18 short papers and 8 industrial papers in this volume were carefully reviewed and selected from 137 submissions. The selected papers covered important topics in the area of service-oriented computing, including foundational issues on service discovery and service-systems design, business process modelling and management, economics of service-systems engineering, as well as services on the cloud, social networks, the Internet of Things (IoT), and data analytics. Learning OpenStack Networking Springer

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.

IBM Redbooks

Programming has become a significant part of connecting theoretical development and scientific application computation. Computer programs and processes that take into account the goals and needs of the user meet with the greatest success, so it behooves software engineers to consider the human element inherent in every line of code they write. Research Anthology on Recent Trends, Tools, and Implications of Computer Programming is a vital reference source that examines the latest scholarly material on trends, techniques, and uses of various programming applications and examines the benefits and challenges of these computational developments. Highlighting a range of topics such as coding standards, software engineering, and computer systems development, this multi-volume book is ideally designed for programmers, computer scientists, software developers, analysts, security experts, IoT software programmers, computer and software engineers, students, professionals, and researchers.

[OpenStack Administration with Ansible](#) IBM Redbooks

This book highlights the importance of security in the design, development and deployment of systems based on Software-Defined Networking (SDN) and Network Functions Virtualization (NFV), together referred to as SDNFV. Presenting a comprehensive guide to the application of security

mechanisms in the context of SDNFV, the content spans fundamental theory, practical solutions, and potential applications in future networks. Topics and features: introduces the key security challenges of SDN, NFV and Cloud Computing, providing a detailed tutorial on NFV security; discusses the issue of trust in SDN/NFV environments, covering roots of trust services, and proposing a technique to evaluate trust by exploiting remote attestation; reviews a range of specific SDNFV security solutions, including a DDoS detection and remediation framework, and a security policy transition framework for SDN; describes the implementation of a virtual home gateway, and a project that combines dynamic security monitoring with big-data analytics to detect network-wide threats; examines the security implications of SDNFV in evolving and future networks, from network-based threats to Industry 4.0 machines, to the security requirements for 5G; investigates security in the Observe, Orient, Decide and Act (OODA) paradigm, and proposes a monitoring solution for a Named Data Networking (NDN) architecture; includes review questions in each chapter, to test the reader's understanding of each of the key concepts described. This informative and practical volume is an essential resource for researchers interested in the potential of SDNFV systems to address a broad range of network security challenges. The work will also be of great benefit to practitioners wishing to design secure next-generation communication networks, or to develop new security-related mechanisms for SDNFV systems.

**Learning OpenStack Networking (Neutron)** "O'Reilly Media, Inc." Get unstuck and start stacking! About This Book Easily fix the nagging problems that commonly plague OpenStack and become the go-to person in your organization Get better equipped to troubleshoot and solve common problems in performance, availability, and automation that confront production-ready OpenStack environments Save time and decrease frustration by solving significant issues that arise from OpenStack deployments pertaining to storage and networking Who This Book Is For You will need a basic understanding of OpenStack, Linux, and Cloud computing. If you have an understanding of Linux, this book will help you leverage that knowledge in the world of OpenStack, giving you confidence to tackle most issues that may arise. What You Will Learn Diagnose and remediate authentication and authorization problems in Keystone Fix common issues with images served through Glance Master the art of troubleshooting Neutron networking Navigate and overcome problems with Nova Troubleshoot and resolve Cinder block storage issues Identify and correct Swift object storage problems Isolate and fix issues caused by Heat orchestration Leverage Ceilometer and other metering and monitoring tools for effective troubleshooting In Detail OpenStack is a collection of software projects that work together to provide a cloud fabric. OpenStack is one of the fastest growing open source projects in history that unlocks cloud computing for everyone. With OpenStack, you are able to create public or private clouds on your own hardware. The flexibility and control afforded by OpenStack puts the cloud within reach of anyone willing to learn this technology. Starting with an introduction to OpenStack troubleshooting tools, we'll walk through each OpenStack service and how you can quickly diagnose, troubleshoot, and correct problems in your OpenStack. Understanding the various projects and how they interact is essential for anyone attempting to troubleshoot an OpenStack cloud. We will start by explaining each of the major components and the dependencies between them, and move on to show you how to identify and utilize an effective set of OpenStack troubleshooting tools and fix common Keystone problems. Next, we will expose you to common errors and problems you may encounter when using the OpenStack Block Storage service (Cinder). We will then examine Heat, the OpenStack Orchestration Service, where you will learn how to trace errors, determine their root cause, and effectively correct the issue. Finally, you will get to know the best practices to architect your OpenStack cloud in order to achieve optimal performance, availability, and reliability. Style and approach This is straight-to-the point guide to fixing your OpenStack cluster. Common problems are identified and suggestions to resolve these problems are presented in a simple, easy-to-understand manner.

### Open Source Software: Mobile Open Source Technologies IBM Redbooks

Learn how you can put the features of OpenStack to work in the real world in this comprehensive path About This Book Harness the abilities of experienced OpenStack administrators and architects, and run your own private cloud successfully Learn how to install, configure, and manage all of the OpenStack core projects including topics on Object Storage, Block Storage, and Neutron Networking services such as LBaaS and FWaaS Get better equipped to troubleshoot and solve common problems in performance, availability, and automation that confront production-ready OpenStack environments Who This Book Is For This course is for those who are new to OpenStack who want to learn the cloud networking fundamentals and get started with OpenStack networking. Basic understanding of Linux Operating System, Virtualization, and Networking, and Storage principles will come in handy. What You Will Learn Get an introduction to OpenStack and its components Store and retrieve data and images using storage components, such as Cinder, Swift, and Glance Install and configure Swift, the OpenStack Object Storage service, including configuring Container Replication between datacenters Gain hands on experience and familiarity with Horizon, the OpenStack Dashboard user interface Learn how to automate OpenStack installations using Ansible and Foreman Follow practical advice and examples for running OpenStack in production Fix common issues with images served through Glance and master

the art of troubleshooting Neutron networking In Detail OpenStack is a collection of software projects that work together to provide a cloud fabric. Learning OpenStack Cloud Computing course is an exquisite guide that you will need to build cloud environments proficiently. This course will help you gain a clearer understanding of OpenStack's components and their interaction with each other to build a cloud environment. The first module, Learning OpenStack, starts with 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. After getting familiar with the fundamentals and application of OpenStack, let's move deeper into the realm of OpenStack. In the second module, OpenStack Cloud Computing Cookbook, preview how to build and operate OpenStack cloud computing, storage, networking, and automation. Dive into Neutron, the OpenStack Networking service, and get your hands dirty with configuring ML2, networks, routers, and distributed virtual routers. Further, you'll learn practical examples of Block Storage, LBaaS, and FBaaS. The final module, Troubleshooting OpenStack, will help you quickly diagnose, troubleshoot, and correct problems in your OpenStack. We will diagnose and remediate issues in Keystone, Glance, Neutron networking, Nova, Cinder block storage, Swift object storage, and issues caused by Heat orchestration. This Learning Path combines some of the best that Packt has to offer in one complete, curated package. It includes content from the following Packt products: Learning OpenStack by Alok Shrivastwa, Sunil Sarat OpenStack Cloud Computing Cookbook - Third Edition by Kevin Jackson, Cody Bunch, Egle Sigler Troubleshooting OpenStack by Tony Campbell Style and approach This course aims to create a smooth learning path that will teach you how to get started with setting up private and public clouds using a free and open source cloud computing platform—OpenStack. Through this comprehensive course, you'll learn OpenStack Cloud computing from scratch to finish and more! Research Anthology on Recent Trends, Tools, and Implications of Computer Programming XML Press 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. OpenStack Operations Guide CRC Press Docs Like CodeLulu.com Preparing for the Certified OpenStack Administrator Exam IBM Redbooks This book constitutes the refereed proceedings of the 10th International IFIP WG 2.13 Conference on Open Source Systems,

OSS 2014, held in San José, Costa Rica, in May 2014. The 16 revised full papers and 16 short papers presented together with 5 poster papers were carefully reviewed and selected from 61 submissions. They have been organized in the following topical sections: open source visualization and reporting; open source in business modeling; open source in mobile and web technologies; open source in education and research; development processes of open source products; testing and assurance of open source projects; and global impact on open source communities and development. The last section consists of five case studies and demonstrations of open source projects.

**Service-Oriented Computing** John Wiley & Sons 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 IBM PowerVC Version 1.3.2 Introduction and Configuration Packt Publishing Ltd

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.

### OpenStack Cloud Security Manning Publications

Due to the increasing availability of affordable internet services, the number of users, and the need for a wider range of multimedia-based applications, internet usage is on the rise. With so many users and such a large amount of data, the requirements of analyzing large data sets leads to the need for further advancements to information processing. Big Data Processing With Hadoop is an essential reference source that discusses possible solutions for millions of users working with a variety of data applications, who expect fast turnaround responses, but encounter issues with processing data at the rate it comes in. Featuring research on topics such as market basket analytics, scheduler load simulator, and writing YARN applications, this book is ideally designed for IoT professionals, students, and engineers seeking coverage on many of the real-world challenges regarding big data.

### Adaptive Resource Management and Scheduling for Cloud Computing John Wiley & Sons

Summary The Spark distributed data processing platform provides an easy-to-implement tool for ingesting, streaming, and processing data from any source. In Spark in Action, Second Edition, you'll learn to take advantage of Spark's core features and incredible processing speed, with applications including real-time computation, delayed evaluation, and machine learning. Spark skills are a hot commodity in enterprises worldwide, and with Spark's powerful and flexible Java APIs, you can reap all the benefits without first learning Scala or Hadoop. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Analyzing enterprise data starts by reading, filtering, and merging files and streams from many sources. The Spark data processing engine handles this varied volume like a champ, delivering speeds 100 times faster than Hadoop systems. Thanks to SQL support, an intuitive interface, and a straightforward multilanguage API, you can use Spark without learning a complex new ecosystem. About the book Spark in Action, Second Edition, teaches you to create end-to-end analytics applications. In this entirely new book, you'll learn from interesting Java-based examples, including a complete data pipeline for processing NASA satellite data. And you'll discover Java, Python, and Scala code samples hosted on GitHub that you can explore and adapt, plus appendices that give you a cheat sheet for installing tools and understanding Spark-specific terms. What's inside Writing Spark

---

applications in Java Spark application architecture Ingestion through files, databases, streaming, and Elasticsearch Querying distributed datasets with Spark SQL About the reader This book does not assume previous experience with Spark, Scala, or Hadoop. About the author Jean-Georges Perrin is an experienced data and software architect. He is France ' s first IBM Champion and has been honored for 12 consecutive years. Table of Contents PART 1 - THE THEORY CRIPPLED BY AWESOME EXAMPLES 1 So, what is Spark, anyway? 2 Architecture and flow 3 The majestic role of the dataframe 4 Fundamentally lazy 5 Building a simple app for deployment 6 Deploying your simple app PART 2 - INGESTION 7 Ingestion from files 8 Ingestion from databases 9 Advanced ingestion: finding data sources and building your own 10 Ingestion through structured streaming PART 3 - TRANSFORMING YOUR DATA 11 Working with SQL 12 Transforming your data 13 Transforming entire documents 14 Extending transformations with user-defined functions 15 Aggregating your data PART 4 - GOING FURTHER 16 Cache and checkpoint: Enhancing Spark ' s performances 17 Exporting data and building full data pipelines 18 Exploring deployment

OpenStack Cloud Application Development IBM Redbooks

Because of the explosion of unstructured data that is generated by individuals and organizations, a new storage paradigm that is called object storage has been developed. Object storage stores data in a flat namespace that scales to trillions of objects. The design of object storage also simplifies how users access data, supporting new types of applications and allowing users to access data by using various methods, including mobile devices and web applications. Data distribution and management are also simplified, allowing greater collaboration across the globe. OpenStack Swift is an emerging open source object storage software platform that is widely used for cloud storage. IBM® Spectrum Scale, which is based on IBM General Parallel File System (IBM GPFS™) technology, is a high-performance and proven product that is used to store data for thousands of mission-critical commercial installations worldwide. Throughout this IBM Redpaper™ publication, IBM Spectrum™ Scale is used to refer to GPFS. The examples in this paper are based on IBM Spectrum Scale™ V4.2.2. IBM Spectrum Scale also automates common storage management tasks, such as tiering and archiving at scale. Together, IBM Spectrum Scale and OpenStack Swift provide an enterprise-class object storage solution that efficiently stores, distributes, and retains critical data. This paper provides instructions about setting up and configuring IBM Spectrum Scale Object Storage that is based on OpenStack Swift. It also provides an initial set of preferred practices that ensure optimal performance and reliability. This paper is intended for administrators who are familiar with IBM Spectrum Scale and OpenStack Swift components.

#### Deploying OpenStack Springer

The distributed computing infrastructure known as ' the Grid ' has undoubtedly been one of the most successful science-oriented large- scale IT projects of the past 20 years. It is now a fully operational international entity, encompassing several hundred computing sites on all continents and giving access to hundreds of thousands of CPU (central processing unit) cores and hundreds of petabytes of storage, all connected by robust national and international scientific networks. It has evolved to become the main computational platform many scientific communities. This book presents lectures from the Enrico Fermi International School of Physics summer school Grid and Cloud computing: Concepts and Practical Applications, held in Varenna, Italy, in July 2014. The school aimed to cover the conceptual and practical aspects of both the Grid and Cloud computing. The proceedings included here are divided into eight chapters, with chapters 1, 2, 3 and 8 covering general applications of Grid and Cloud computing in various scientific fields, while chapters 4, 5, 6 and 7 discuss specific technical areas of Grid and Cloud structures. The book will be of interest to all those whose work involves the use of the Grid or Cloud computing.