

How To Install Mariadb Galera Cluster On Ubuntu 16 04

Today's system administrators deal with a vast number of situations, operating systems, software packages, and problems. Those who are in the know have kept their copy of Linux Server Hacks close at hand to ease their burden. And while this helps, it's not enough: any sys admin knows there are many more hacks, cool tips, and ways of solving problems than can fit in a single volume (one that mere mortals can lift, that is). Which is why we created Linux Server Hacks, Volume Two, a second collection of incredibly useful tips and tricks for finding and using dozens of open source tools you can apply to solve your sys admin problems. The power and flexibility of Linux and Open Source means that there is an astounding amount of great software out there waiting to be applied to your sys admin problems -- if only you knew about it and had enough information to get started. Hence, Linux Server Hacks, Volume Two. This handy reference offers 100 completely new server management tips and techniques designed to improve your productivity and sharpen your administrative skills. Each hack represents a clever way to accomplish a specific task, saving you countless hours of searching for the right answer. No more sifting through man pages, HOWTO websites, or source code comments -- the only resource you need is right here. And you don't have to be a system administrator with hundreds of boxen to get something useful from this book as many of the hacks apply equally well to a single system or a home network. Compiled by experts, these hacks not only give you the step-by-step instructions necessary to implement the software, but they also provide the context to truly enable you to learn the technology. Topics include: Authentication Remote GUI connectivity Storage management File sharing and synchronizing resources Security/lockdown instruction Log files and monitoring Troubleshooting System rescue, recovery, and repair Whether they help you recover lost data, collect information from distributed clients, or synchronize administrative environments, the solutions found in Linux Server Hacks, Volume Two will simplify your life as a system administrator.

This IBM® Redpaper™ publication shows you how to deploy a database instance within a container using an IBM Cloud™ Private cluster on IBM Z®. A preinstalled IBM Spectrum™ Scale 5.0.3 cluster file system provides back-end storage for the persistent volumes bound to the database. A container is a standard unit of software that packages code and all its dependencies, so the application runs quickly and reliably from one computing environment to another. By default, containers are ephemeral. However, stateful applications, such as databases, require some type of persistent storage that can survive service restarts or container crashes. IBM provides several products helping organizations build an environment on an IBM Z infrastructure to develop and manage containerized applications, including dynamic provisioning of persistent

volumes. As an example for a stateful application, this paper describes how to deploy the relational database MariaDB using a Helm chart. The IBM Spectrum Scale V5.0.3 cluster file system is providing back-end storage for the persistent volumes. This document provides step-by-step guidance regarding how to install and configure the following components: IBM Cloud Private 3.1.2 (including Kubernetes) Docker 18.03.1-ce IBM Storage Enabler for Containers 2.0.0 and 2.1.0 This Redpaper demonstrates how we set up the example for a stateful application in our lab. The paper gives you insights about planning for your implementation. IBM Z server hardware, the IBM Z hypervisor z/VM®, and the IBM Spectrum Scale cluster file system are prerequisites to set up the example environment. The Redpaper is written with the assumption that you have familiarity with and basic knowledge of the software products used in setting up the environment. The intended audience includes the following roles: Storage administrators IT/Cloud administrators Technologists IT specialists

"With an easy, step-by-step approach, this guide shows beginners how to install, use, and maintain the world's most popular open source database: MySQL. You'll learn through real-world examples and many practical tips, including information on how to improve database performance. Database systems such as MySQL help data handling for organizations large and small handle data, providing robust and efficient access in ways not offered by spreadsheets and other types of data stores. This book is also useful for web developers and programmers interested in adding MySQL to their skill sets. Topics include: Installation and basic administration ; Introduction to databases and SQL ; Functions, subqueries, and other query enhancements ; Improving database performance ; Accessing MySQL from popular languages" --

A guide for MySQL administrators covers such topics as benchmarking, server performance, indexing, queries, hardware optimization, replication, scaling, cloud hosting, and backup and recovery.

Learn how to transition from Excel-based business intelligence (BI) analysis to enterprise stacks of open-source BI tools. Select and implement the best free and freemium open-source BI tools for your company's needs and design, implement, and integrate BI automation across the full stack using agile methodologies. Business Intelligence Tools for Small Companies provides hands-on demonstrations of open-source tools suitable for the BI requirements of small businesses. The authors draw on their deep experience as BI consultants, developers, and administrators to guide you through the extract-transform-load/data warehousing (ETL/DWH) sequence of extracting data from an enterprise resource planning (ERP) database freely available on the Internet, transforming the data, manipulating them, and loading them into a relational database. The authors demonstrate how to extract, report, and dashboard key performance indicators (KPIs) in a visually appealing format from the relational database management system (RDBMS). They model the selection and implementation of free and freemium tools such as Pentaho Data Integrator and

Talend for ELT, Oracle XE and MySQL/MariaDB for RDBMS, and QlikSense, Power BI, and MicroStrategy Desktop for reporting. This richly illustrated guide models the deployment of a small company BI stack on an inexpensive cloud platform such as AWS. What You'll Learn You will learn how to manage, integrate, and automate the processes of BI by selecting and implementing tools to: Implement and manage the business intelligence/data warehousing (BI/DWH) infrastructure Extract data from any enterprise resource planning (ERP) tool Process and integrate BI data using open-source extract-transform-load (ETL) tools Query, report, and analyze BI data using open-source visualization and dashboard tools Use a MOLAP tool to define next year's budget, integrating real data with target scenarios Deploy BI solutions and big data experiments inexpensively on cloud platforms Who This Book Is For Engineers, DBAs, analysts, consultants, and managers at small companies with limited resources but whose BI requirements have outgrown the limitations of Excel spreadsheets; personnel in mid-sized companies with established BI systems who are exploring technological updates and more cost-efficient solutions

MariaDB is a database server that offers drop-in replacement functionality for MySQL. Built by some of the original authors of MySQL, with assistance from the broader community of free and open source software developers, MariaDB offers a rich set of feature enhancements to MySQL, including alternate storage engines, server optimizations, and patches. MariaDB Crash Course teaches you all you need to know to be immediately productive with MariaDB. Master trainer Ben Forta introduces all the essentials through a series of quick, easy-to-follow, hands-on lessons. Instead of belaboring database theory and relational design, Forta focuses on teaching solutions for the majority of users who simply want to interact with data. Learn how to: Retrieve and sort data Filter data using comparisons, regular expressions, and full text search Join relational data Create and alter tables Insert, update, and delete data Leverage the power of stored procedures and triggers Use views and cursors Manage transactional processing Create user accounts and manage security via access control

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.

Mitigate the risks involved in migrating away from a proprietary database platform toward MariaDB's open source database engine. This book will help you assess the risks and the work involved, and ensure a successful migration. Migrating to MariaDB describes the process and lessons learned during a migration from a proprietary database management engine to the MariaDB open source solution. The book discusses the drivers for making the decision and change, walking you through all aspects of the process from evaluating the licensing, navigating the pitfalls and hurdles of a migration, through to final implementation on the new platform. The book highlights the cost-effectiveness of MariaDB and how the licensing worries are simplified in comparison to running on a proprietary platform. You'll learn to do your own risk assessment, to identify database and application code that may need to be modified or re-implemented, and to identify MariaDB features to provide the security and failover protection needed by corporate customers. Let the author's experience in migrating a financial firm to MariaDB inform your own efforts, helping you to develop a road map for both

technical and political success within your own organization as you migrate away from proprietary lock-in toward MariaDB's open source solution. What You'll Learn Evaluate and compare licensing costs between proprietary databases and MariaDB Perform a proper risk assessment to inform your planning and execution of the migration Build a migration road map from the book's example that is specific to your situation Make needed application changes and migrate data to the MariaDB open source database engine Who This Book Is For Technical professionals (including database administrators, programmers, and technical management) who are interested in migrating away from a proprietary database platform toward MariaDB's open source database engine and need to assess the risks and the work involved

ZFS improves everything about systems administration. Once you peek under the hood, though, ZFS' bewildering array of knobs and tunables can overwhelm anyone. ZFS experts can make their servers zing—and now you can, too, with *FreeBSD Mastery: Advanced ZFS*. This small book teaches you to: •Use boot environments to make the riskiest sysadmin tasks boring •Delegate filesystem privileges to users •Containerize ZFS datasets with jails •Quickly and efficiently replicate data between machines •split layers off of mirrors •optimize ZFS block storage •handle large storage arrays •select caching strategies to improve performance •manage next-generation storage hardware •identify and remove bottlenecks •build screaming fast database storage •dive deep into pools, metaslabs, and more! Whether you manage a single small server or international datacenters, simplify your storage with *FreeBSD Mastery: Advanced ZFS*.

Achieve enterprise automation in your Linux environment with this comprehensive guide **Key Features** Automate your Linux infrastructure with the help of practical use cases and real-world scenarios Learn to plan, build, manage, and customize OS releases in your environment Enhance the scalability and efficiency of your infrastructure with advanced Linux system administration concepts **Book Description** Automation is paramount if you want to run Linux in your enterprise effectively. It helps you minimize costs by reducing manual operations, ensuring compliance across data centers, and accelerating deployments for your cloud infrastructures. Complete with detailed explanations, practical examples, and self-assessment questions, this book will teach you how to manage your Linux estate and leverage Ansible to achieve effective levels of automation. You'll learn important concepts on standard operating environments that lend themselves to automation, and then build on this knowledge by applying Ansible to achieve standardization throughout your Linux environments. By the end of this Linux automation book, you'll be able to build, deploy, and manage an entire estate of Linux servers with higher reliability and lower overheads than ever before. What you will learn Perform large-scale automation of Linux environments in an enterprise Overcome the common challenges and pitfalls of extensive automation Define the business processes needed to support a large-scale Linux environment Get well-versed with the most effective and reliable

patch management strategies Automate a range of tasks from simple user account changes to complex security policy enforcement Learn best practices and procedures to make your Linux environment automatable Who this book is for This book is for anyone who has a Linux environment to design, implement, and maintain. Open source professionals including infrastructure architects and system administrators will find this book useful. You're expected to have experience in implementing and maintaining Linux servers along with knowledge of building, patching, and maintaining server infrastructure. Although not necessary, knowledge of Ansible or other automation technologies will be beneficial.

This IBM® Redbooks® publication is Volume 4 of a series of books entitled The Virtualization Cookbook for IBM z Systems. The other volumes in the series are: The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3, SG24-8147 The Virtualization Cookbook for IBM z Systems Volume 2: Red Hat Enterprise Linux 7.1 Servers, SG24-8303 The Virtualization Cookbook for IBM z Systems Volume 3: SUSE Linux Enterprise Server 12, SG24-8890 It is advised that you start with Volume 1 of this series, because the IBM z/VM® Hypervisor is the foundation for installing Linux on IBM z™ Systems.

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.

Python is a powerful, expressive programming language that's easy to learn and fun to use! But books about learning to program in Python can be kind of dull, gray, and boring, and that's no fun for anyone. Python for Kids brings Python to life and brings you (and your parents) into the world of programming. The ever-patient Jason R. Briggs will guide you through the basics as you experiment with unique (and often hilarious) example programs that feature ravenous monsters, secret agents, thieving ravens, and more. New terms are defined; code is colored, dissected, and explained; and quirky, full-color illustrations keep things on the lighter side. Chapters end with programming puzzles designed to stretch your brain and strengthen your understanding. By the end of the book you'll have programmed two complete games: a clone of the famous Pong and "Mr. Stick Man Races for the Exit"—a platform game with jumps, animation, and much more. As you strike out on your programming adventure, you'll learn how to:

- Use fundamental data structures like lists, tuples, and maps
- Organize and reuse your code with functions and modules
- Use control structures like loops and conditional statements
- Draw shapes and patterns with Python's turtle module
- Create games, animations, and other graphical wonders with tkinter

Why should serious adults have all the fun? Python for Kids is your ticket into the amazing world of computer programming. For kids ages 10+ (and their parents) The code in this book runs on almost anything: Windows, Mac, Linux, even an OLPC laptop or Raspberry Pi!

SELinux: Bring World-Class Security to Any Linux Environment! SELinux offers Linux/UNIX integrators, administrators, and developers a state-of-the-art platform for building and maintaining highly secure solutions. Now that SELinux is included in the Linux 2.6 kernel—and delivered by default in Fedora Core, Red Hat Enterprise Linux, and other major distributions—it's easier than ever to take advantage of its benefits. SELinux by Example is the first complete, hands-on guide to using SELinux in production environments. Authored by three leading SELinux researchers and developers, it illuminates every facet of working with SELinux, from its architecture and security object model to its policy language. The book thoroughly explains SELinux sample policies—including the powerful new Reference Policy—showing how to quickly adapt them to your unique environment. It also contains a comprehensive SELinux policy language

reference and covers exciting new features in Fedora Core 5 and the upcoming Red Hat Enterprise Linux version 5. • Thoroughly understand SELinux's access control and security mechanisms • Use SELinux to construct secure systems from the ground up • Gain fine-grained control over kernel resources • Write policy statements for type enforcement, roles, users, and constraints • Use optional multilevel security to enforce information classification and manage users with diverse clearances • Create conditional policies that can be changed on-the-fly • Define, manage, and maintain SELinux security policies • Develop and write new SELinux security policy modules • Leverage emerging SELinux technologies to gain even greater flexibility • Effectively administer any SELinux system

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

into the design of the Kubernetes platform, the reader will be exposed to the advanced features and best practices of Kubernetes. This book will be an advanced level book which will provide a pathway to master Kubernetes

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! 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
OpenStack Cloud Computing CookbookPackt Publishing Ltd
Orchestrate and automate your OpenStack cloud operator tasks with Ansible 2.0
About This Book Automate real-world OpenStack cloud operator administrative
tasks Construct a collection of the latest automation code to save time on
managing your OpenStack cloud Manage containers on your cloud and check
the health of your cloud using Nagios Who This Book Is For This book is aimed
at OpenStack-based cloud operators and infrastructure and sys administrators
who have some knowledge of OpenStack and are seeking to automate taxing
and manual tasks. This book is also for people new to automating cloud
operations in general and the DevOps practice in particular. What You Will Learn
Efficiently execute OpenStack administrative tasks Familiarize yourself with how
Ansible 2 works and assess the defined best practices Create Ansible 2
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 current production-ready projects
Explore OpenStack CLI tools and learn 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. We
start with a brief overview of OpenStack and Ansible 2 and highlight some best
practices. Each chapter will provide an introduction to handling various Cloud
Operator administration tasks such as managing containers within your cloud;
setting up/utilizing open source packages for monitoring; creating multiple
users/tenants; taking instance snapshots; and customizing your cloud to run
multiple active regions. Each chapter will also supply a step-by-step tutorial on
how to automate these tasks with Ansible 2. 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 on the market today. Style and approach This book is a
concise, fast-paced guide filled with real-world scenarios that will execute
OpenStack administrative tasks efficiently. It serves as a quick reference guide
for not just OpenStack functions, but also for creating future Ansible code.
Get a comprehensive overview on how to set up and design an effective
database with MySQL. This thoroughly updated edition covers MySQL's latest
version, including its most important aspects. Whether you're deploying an
environment, troubleshooting an issue, or engaging in disaster recovery, this
practical guide provides the insights and tools necessary to take full advantage of
this powerful RDBMS. Authors Vinicius Grippa and Sergey Kuzmichev from
Percona show developers and DBAs methods for minimizing costs and
maximizing availability and performance. You'll learn how to perform basic and

advanced querying, monitoring and troubleshooting, database management and security, backup and recovery, and tuning for improved efficiency. This edition includes new chapters on high availability, load balancing, and using MySQL in the cloud. Get started with MySQL and learn how to use it in production Deploy MySQL databases on bare metal, on virtual machines, and in the cloud Design database infrastructures Code highly efficient queries Monitor and troubleshoot MySQL databases Execute efficient backup and restore operations Optimize database costs in the cloud Understand database concepts, especially those pertaining to MySQL

The Essential Guide to Scaling Your Business Solution Written by Oracle ACE Director and MySQL expert Ronald Bradford, with coauthor Chris Schneider, *Effective MySQL: Replication Techniques in Depth* describes what is needed to understand and implement MySQL replication to build scalable solutions. This book includes detailed syntax examples to demonstrate the features, options, and limitations of native MySQL replication. Providing an evaluation of various new replication features and additional third-party product implementations, this Oracle Press guide helps to ensure your MySQL environment can support the various high-availability needs of your business. Master the strengths and limitations of native asynchronous replication in a MySQL topology Identify the important features to improve replication for growing business requirements Recognize the key business factors to determine your optimal highavailability needs Understand the benefits of using MySQL replication for failover scenarios Identify the key configuration variables and SQL commands affecting master/slave replication Learn about the advancements in replication techniques provided by new products, including Tungsten Replicator and Galera Optimize your replication management with various utilities and toolkits Find additional detailed information and presentations at EffectiveMySQL.com.

Create high availability clusters to enhance system performance using CentOS 7 About This Book Master the concepts of high performance and high availability to eliminate performance bottlenecks Maximize the uptime of services running in a CentOS 7 cluster A step-by-step guide that will provide knowledge of methods and approaches to optimize the performance of CentOS clusters Who This Book Is For This book is targeted at system administrators: those who want a detailed, step-by-step guide to learn how to set up a high-availability CentOS 7 cluster, and those who are looking for a reference book to help them learn or refresh the necessary skills to ensure their systems and respective resources are utilized optimally. No previous knowledge of high-availability systems is needed, though the reader is expected to have at least some degree of familiarity with any spin-off of the Fedora family of Linux distributions, preferably CentOS. What You Will Learn Install a CentOS 7 cluster and network infrastructure Configure firewall, networking, and clustering services and settings Set up and test a HAC (high-availability cluster) to host an Apache web server and a MariaDB database server Monitor performance and availability Identify bottlenecks and troubleshoot

issues Improve performance and ensure high availability In Detail CentOS is the enterprise level Linux OS, which is 100% binary compatible to Red Hat Enterprise Linux (RHEL). It acts as a free alternative to RedHat's commercial Linux offering, with only a change in the branding. A high performance cluster consists in a group of computers that work together as one set parallel, hence minimizing or eliminating the downtime of critical services and enhancing the performance of the application. Starting with the basic principles of clustering, you will learn the necessary steps to install a cluster with two CentOS 7 servers. We will then set up and configure the basic required network infrastructure and clustering services. Further, you will learn how to take a proactive approach to the split-brain issue by configuring the failover and fencing of the cluster as a whole and the quorum of each node individually. Further, we will be setting up HAC and HPC clusters as a web server and a database server. You will also master the art of monitoring performance and availability, identifying bottlenecks, and exploring troubleshooting techniques. At the end of the book, you'll review performance-tuning techniques for the recently installed cluster, test performance using a payload simulation, and learn the necessary skills to ensure that the systems, and the corresponding resources and services, are being utilized to their best capacity. Style and approach An easy-to-follow and step-by-step guide with hands-on instructions to set up real-world simple cluster scenarios that will start you on the path to building more complex applications on your own.

This book is intended for intermediate users who want to learn how to administrate a MariaDB server or a set of servers. It is aimed at MariaDB users, and hence working knowledge of MariaDB is a prerequisite.

This IBM® Redpaper publication describes IBM Spectrum® LSF® Suite best practices installation topics, application checks for workload management, and high availability configurations by using theoretical knowledge and hands-on exercises. These findings are documented by way of sample scenarios. This publication addresses topics for sellers, IT architects, IT specialists, and anyone who wants to implement and manage a high-performing workload management solution with LSF. Moreover, this guide provides documentation to transfer how-to-skills to the technical teams, and solution guidance to the sales team. This publication compliments documentation that is available at IBM Knowledge Center, and aligns with educational materials that are provided by IBM Systems. 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.

Cacti adalah salah satu sistem monitoring jaringan yang bersifat open. Banyak penggunaan aktif yang berpartisipasi di ruang Forum Cacti. Dan hebatnya lagi banyak perusahaan kecil dan sampai kelas ISP besar menggunakan sistem monitoring ini. Buku ini membahas bagaimana memulai membangun Cacti menggunakan sistem operasi Windows dan Linux. Buku ini berusaha menjelaskan dari sisi praktis dan contoh real penggunaannya di lapangan/tempat kerja. Dilengkapi juga dengan ilustrasi gambar "step-by-step" untuk memudahkan pembaca.

60 Menit Belajar Python menyampaikan dengan sesederhana mungkin bagaimana memulai Python disertai contoh langkah demi langkahnya demi untuk memudahkan pemahaman. Tak lupa juga memberikan tambahan tip/cara yang semoga dapat membantu. Untuk lebih mempermudah pembaca, juga telah diperkaya dengan contoh program singkat dan template program. Diharapkan akan membantu banyak dalam proses (mempercepat) belajar dan pada akhirnya

dapat dimanfaatkan untuk keperluan yang dibutuhkan.

This book is aimed at system administrators/architects or DBAs who want to learn more about how to grow their current infrastructure to support larger traffic. Before beginning with this book, we expect you to be well-practiced with MySQL/MariaDB for common usage. You will be able to get a grasp quickly if you are comfortable with learning and building large infrastructures for MariaDB using Linux.

Over 110 effective recipes to help you build and operate OpenStack cloud computing, storage, networking, and automation About This Book Explore many new features of OpenStack's Juno and Kilo releases Install, configure, and administer core projects with the help of OpenStack Object Storage, Block Storage, and Neutron Networking services Harness the abilities of experienced OpenStack administrators and architects, and run your own private cloud successfully Practical, real-world examples of each service and an accompanying Vagrant environment that helps you learn quickly In Detail OpenStack Open Source software is one of the most used cloud infrastructures to support software development and big data analysis. It is developed by a thriving community of individual developers from around the globe and backed by most of the leading players in the cloud space today. It is simple to implement, massively scalable, and can store a large pool of data and networking resources. OpenStack has a strong ecosystem that helps you provision your cloud storage needs. Add OpenStack's enterprise features to reduce the cost of your business. This book will show you the steps to build up a private cloud environment. At the beginning, you'll discover the uses of cloud services such as the identity service, image service, and compute service. You'll dive into Neutron, the OpenStack Networking service, and get your hands dirty with configuring ML2, networks, routers, and Distributed Virtual Routers. You'll then gather more expert knowledge on OpenStack cloud computing by managing your cloud's security and migration. After that, we delve in to OpenStack Object storage and how to manage servers and work with objects, cluster, and storage functionalities. Also, as you go deeper into the realm of OpenStack, you'll learn practical examples of Block storage, LBaaS, and FWaaS: installation and configuration covered ground up. Finally, you will learn OpenStack dashboard, Ansible and Foreman, Keystone, and other interesting topics. What You Will Learn Understand, install, configure, and manage Nova—the OpenStack Cloud Compute resource Configure ML2, networks, routers, and Distributed Virtual Routers with Neutron Use and secure Keystone, the OpenStack Authentication service Install and set up Swift and Container Replication between datacenters Gain hands-on experience and familiarity with Horizon, the OpenStack Dashboard user interface Automate complete solutions with our recipes on Heat, the OpenStack Orchestration service Use Ansible and Foreman to automate OpenStack installations successfully Follow practical advice and examples to run OpenStack in production Who This Book Is For This book is aimed at cloud system engineers,

Jaeger Extend Kubernetes working with Kubernetes API, plugins, and webhooks
Who this book is for If you are a system administrator or a cloud developer with working knowledge of Kubernetes and are keen to master its advanced features, along with learning everything from building microservices to utilizing service meshes, Mastering Kubernetes is for you. Basic familiarity with networking concepts will be helpful.

Learning cryptography and security is fun instead of saying it hard or complex. This book is written in cookbook style and covers all the major crypto function with the sample code using the major python crypto libraray like (cryptography/pycrypto/jwcrypto), which will come handy for python crypto developers from beginner to advanced in their daily use.

Quickly get up to speed with MariaDB—the leading, drop-in replacement for MySQL, through this practical tutorial
About This Book Get to know the basic SQL queries so you can quickly start using MariaDB
Take control of your data through the advanced features of MariaDB
Exploit the full potential of MariaDB's exclusive features through quick, practical examples
Who This Book Is For If you don't know the SQL language, but you want to quickly jump into the SQL world and learn how to use MariaDB, or if you already know how to use MySQL but you want to go further, then this book is ideal for you.
What You Will Learn
Install and configure MariaDB
Create databases, tables, and indexes
Import and export data from and to external files
Work with views and virtual columns
Create, read, update, and delete records in your database
Use dynamic columns
Set up a powerful full-text search system
Access your external data from MariaDB through the CONNECT engine
In Detail This book will take you through all the nitty-gritty parts of MariaDB, right from the creation of your database all the way to using MariaDB's advanced features. At the very beginning, we show you the basics, that is, how to install MariaDB. Then, we walk you through the databases and tables of MariaDB, and introduce SQL in MariaDB. You will learn about all the features that have been added in MariaDB but are absent in MySQL. Moving on, you'll learn to import and export data, views, virtual columns, and dynamic columns in MariaDB. Then, you'll get to grips with full-text searches and queries in MariaDb. You'll also be familiarized with the CONNECT storage engine. At the end of the book, you'll be introduced to the community of MariaDB.
Style and approach This is a complete guide that uses concrete examples to help you understand and exploit the full potential of MariaDB.

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.

The Cisco expert guide to planning, deploying, and operating virtual routing with the CSR 1000V Cloud Services Router
Virtual routing and the Cisco Cloud Services Router (CSR 1000V) are key enablers of today's revolutionary shift to elastic cloud applications and low-cost virtualized networking. Now, there's an authoritative, complete guide to building real solutions with the Cisco CSR 1000V platform. Three leading experts cover every essential building block, present key use cases and configuration examples, illuminate design and deployment scenarios, and show how the CSR 1000V platform and APIs can enable state-of-the-art software-defined networks (SDN). Drawing on extensive early adopter experience, they illuminate crucial OS and hypervisor details, help you overcome migration challenges, and offer practical guidance for monitoring and operations. This guide is an essential resource for all technical professionals planning or deploying data center and enterprise cloud services, and for all cloud network operators utilizing the Cisco CSR 1000V or future Cisco virtual routing platforms. - Review the fundamentals of cloud virtualization, multitenant data-center design,

and software-defined networking · Understand the Cisco CSR 1000V's role, features, and infrastructure requirements · Compare server hypervisor technologies for managing VM hardware with CSR 1000V deployments · Understand CSR 1000V software architecture, control and data-plane design, licensing requirements, and packet flow · Walk through common virtual router scenarios and configurations, including multiple cloud and data center examples · Integrate CSR 1000V into the OpenStack SDN framework, and use its APIs to solve specific problems · Master a best-practice workflow for deploying the CSR 1000V · Use the Cisco management tools to automate, orchestrate, and troubleshoot virtualized routing

Category: Networking/Cloud Computing Covers: Cloud Services Router This book is part of the Networking Technology Series from Cisco Press, which offers networking professionals valuable information for constructing efficient networks, understanding new technologies, and building successful careers

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.

This is the latest edition of the book that application developers worldwide have used to master MySQL]€]now updated for MySQL 8 and beyond. As you would expect, this book shows how

to code all the essential SQL statements for working with a MySQL database. You'll use these statements every day to have MySQL do more of your work for you. But beyond that, it shows how to work with classic MySQL features that take you to new level, such as summary queries, subqueries, functions, views, transactions, stored procedures, triggers, and security. It shows how to take advantage of newer MySQL features such as window functions, Common Table Expressions (CTE), and roles for database security. It shows how to design a database, including how to use MySQL Workbench to create and implement the design. It even presents a starting set of skills for a database administrator (DBA) if you're interested in that career path or if you need to be your own DBA. In short, it's a must-have guide for anyone who works with MySQL, beginning and experienced developers alike.

A practical cookbook, filled with advanced recipes, and plenty of code and commands used for illustration, which will make your learning curve easy and quick. This book is for anyone who wants to learn more about databases in general or MariaDB in particular. Some familiarity with SQL databases is assumed, but the recipes are approachable to almost anyone with basic database skills.

[Copyright: 479da95855714176f291540382459002](#)