

## Ubuntu 16 04 LTS Server: Administration And Reference

This book constitutes the refereed conference proceedings of the 31st British International Conference on Databases, BICOD 2017 - formerly known as BNCOD (British National Conference on Databases) - held in London, UK, in July 2017. The 17 revised full papers were carefully reviewed and selected from numerous submissions. The papers cover a wide range of topics such as data cleansing, data integration, data wrangling, data mining and knowledge discovery, graph data and knowledge graphs, intelligent data analysis, approximate and flexible querying, data provenance and ontology-based data access. They are organized in the following topical sections: data wrangling and data integration; data analysis and data mining; graph data querying and analysis; multidimensional data and data quality; and distributed and multimedia data management.

This book is designed as an Ubuntu 16.04 LTS server administration and quick reference. The purpose of this book is to help the Ubuntu users in doing daily tasks weather it is basic or advanced tasks. The book contains lot of screen shots and output of commands to accomplice the required task. This book emphasis on practical aspect of Ubuntu administration therefore you will find procedures with required theory only. This book cover LVM, ftp, user administration, firewall, NFS, logs and software management etc.

Follow this handbook to build, configure, tune, and secure Apache Cassandra databases. Start with the installation of Cassandra and move on to the creation of a single instance, and then a cluster of Cassandra databases. Cassandra is increasingly a key player in many big data environments, and this book shows you how to use Cassandra with Apache Spark, a popular big data processing framework. Also covered are day-to-day topics of importance such as the backup and recovery of Cassandra databases, using the right compression and compaction strategies, and loading and unloading data. Expert Apache Cassandra Administration provides numerous step-by-step examples starting with the basics of a Cassandra database, and going all the way through backup and recovery, performance optimization, and monitoring and securing the data. The book serves as an authoritative and comprehensive guide to the building and management of simple to complex Cassandra databases. The book: Takes you through building a Cassandra database from installation of the software and creation of a single database, through to complex clusters and data centers Provides numerous examples of actual commands in a real-life Cassandra environment that show how to confidently configure, manage, troubleshoot, and tune Cassandra databases Shows how to use the Cassandra configuration properties to build a highly stable, available, and secure Cassandra database that always operates at peak efficiency What You'll Learn Install the Cassandra software and create your first database Understand the Cassandra data model, and the internal architecture of a Cassandra database Create your own Cassandra cluster, step-by-step Run a Cassandra cluster on Docker Work with Apache Spark by connecting to a Cassandra database Deploy Cassandra clusters in your data center, or on Amazon EC2 instances Back up and restore mission-critical Cassandra

databases Monitor, troubleshoot, and tune production Cassandra databases, and cut your spending on resources such as memory, servers, and storage Who This Book Is For Database administrators, developers, and architects who are looking for an authoritative and comprehensive single volume for all their Cassandra administration needs. Also for administrators who are tasked with setting up and maintaining highly reliable and high-performing Cassandra databases. An excellent choice for big data administrators, database administrators, architects, and developers who use Cassandra as their key data store, to support high volume online transactions, or as a decentralized, elastic data store.

Want to learn a new skill? Expand the technology that you work with? This is the latest version of Neil's books on the Ubuntu operating system and revised to tech the newest version of Ubuntu (16.04 LTS). This book covers the basics of understanding how to use the latest long term support edition of Ubuntu (16.04 LTS). In this book we use Ubuntu 16.04 LTS to learn multiple fundamentals in using Linux and later will go through the process of creating a web server. 80% of websites are driven by Linux servers.

Understanding the basics and expanding upon this will provide great career opportunities and a great skill as well. We start simple and the reader does not need any prior knowledge. We will make baby steps and slowly work ourselves up to configuring the Ubuntu Server to be a functional web server. There will still be much to learn, but within a few hours you can have your own Linux server setup, understand the basics, and also have WordPress loaded into it. We cover installing packages, creating files in nano, LAMP stack (Linux Apache MySQL PHP), and try to do so in a practical way so that you can finish this guide with something to show off. This book covers multiple ways to get started, including two different VPS providers, installing the OS onto a computer, and configuration a virtual machine using Oracle's VirtualBox.

Administration and Quick Reference

Smart and Innovative Trends in Next Generation Computing Technologies

Learn to Build Systems for Your Business Using Free and Open Source Software

SQL Server on Linux

Third International Conference, NGCT 2017, Dehradun, India, October 30-31, 2017, Revised Selected Papers, Part I

Data Analytics

Bring the performance and security of SQL Server to Linux About This Book Design and administer your SQL Server solution on the open source Linux platform Install, configure, and fine-tune your database application for maximum performance An easy-to-follow guide teaching you how to implement various SQL Server CTP 2.x offerings on Linux—from installation to administration Who This Book Is For This book is for the Linux users who want to learn SQL Server on their favorite Linux distributions. It is not important if you are experienced database user or a beginner as we are starting from scratch. However, it is recommended that you have basic knowledge about relational models. More advanced readers can pick the chapters of their interest and study specific topics immediately. Users from Windows platform can also benefit

from this book to expand their frontiers and become equally efficient on both platforms. What You Will Learn Install and set up SQL Server CTP 2.x on Linux Create and work with database objects using SQL Server on Linux Configure and administer SQL Server on Linux-based systems Create and restore database back-ups Protect sensitive data using the built-in cryptographic features Optimize query execution using indexes Improve query execution time by more than 10x using in-memory OLTP Track row-versioning using temporal tables In Detail Microsoft's launch of SQL Server on Linux has made SQL Server a truly versatile platform across different operating systems and data-types, both on-premise and on-cloud. This book is your handy guide to setting up and implementing your SQL Server solution on the open source Linux platform. You will start by understanding how SQL Server can be installed on supported and unsupported Linux distributions. Then you will brush up your SQL Server skills by creating and querying database objects and implementing basic administration tasks to support business continuity, including security and performance optimization. This book will also take you beyond the basics and highlight some advanced topics such as in-memory OLTP and temporal tables. By the end of this book, you will be able to recognize and utilize the full potential of setting up an efficient SQL Server database solution in your Linux environment. Style and approach This book follows a step-by-step approach to teach readers the concepts of SQL Server on Linux using the bash command line and SQL programming language through examples which can easily be adapted and applied in your own solutions.

This book gathers selected high-quality research papers presented at the Sixth International Congress on Information and Communication Technology, held at Brunel University, London, on February 25 – 26, 2021. It discusses emerging topics pertaining to information and communication technology (ICT) for managerial applications, e-governance, e-agriculture, e-education and computing technologies, the Internet of things (IoT) and e-mining. Written by respected experts and researchers working on ICT, the book offers a valuable asset for young researchers involved in advanced studies. The book is presented in four volumes.

Explore and learn the key building blocks of Microsoft Azure services and tools for implementing a disaster-recovery solution of any on-premises or cloud-based application. In this book, you will go through various aspects of planning, designing, and configuring a robust recovery solution on Azure. Introducing Disaster Recovery with Microsoft Azure starts by explaining the disaster-recovery landscape and how Azure disaster recovery is different from the traditional approach. You will learn how to leverage Azure site recovery and various Azure-based services to design and implement a recovery solution and much more. Moving forward, you will design and implement various scenarios such as on-premises to Azure, Azure to Azure, and on-premises to on-premises disaster recovery. You will also learn common considerations and technicalities of implementing recovery solutions for various multi-tier, monolithic, and modern micro-services enterprise applications. Finally, you will go through real-life examples, scenarios, and exercises. After reading this book, you will be able to design and implement disaster recovery on Azure in different scenarios. You will also look at a few real-world scenarios that will provide more practical insights. What You Will Learn Discover the fundamental building blocks of disaster recovery on Azure Examine various application-specific considerations for disaster recovery Leverage various PaaS capabilities to achieve maximum benefit Design and implement a multi-regional Azure to Azure disaster recovery plan Who This Book Is For Consultants, architects, and Azure administrators.

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.

16th International Symposium, I-SPAN 2019, Naples, Italy, September 16-20, 2019, Proceedings

Ansible by Examples

Proceedings of Sixth International Congress on Information and Communication Technology

100+ Automation Examples to Automate Security and Verify Compliance for IT Modern Infrastructure

Over 30 recipes for implementing deep neural networks in Python

Mastering Ubuntu 16.04 Lts Server Installation and Administration

Understanding Services and Tools for Implementing a Recovery Solution

These proceedings gather papers presented at the Cyber Security Education Stream and Cyber Security Technology Stream of The National Cyber Summit's Research Track, and report on the latest advances in areas ranging from software security to cyber attack detection and modeling; the use of machine learning in cyber security; legislation and policy; surveying small businesses; cyber competition, and so on. Understanding the latest capabilities in cyber security is the best way to prepare users and organizations for potential negative events. Consequently, this book will be of interest to cyber security researchers, educators and practitioners, as well as students who want to learn about cyber security.

Leverage the power of deep learning and Keras to develop smarter and more efficient data models Key Features Understand different neural networks and their implementation using Keras Explore recipes for training and fine-tuning your neural network models Put your deep learning knowledge to practice with real-world use-cases, tips, and tricks Book Description Keras has quickly emerged as a popular deep learning library. Written in Python, it allows you to train convolutional as well as recurrent neural networks with speed and accuracy. The Keras Deep Learning Cookbook shows you how to tackle different problems encountered while training efficient deep learning models, with the help of the popular Keras library. Starting with installing and setting up Keras, the book demonstrates how you can perform deep learning with Keras on the TensorFlow. From loading data to fitting and evaluating your model for optimal performance, you will work through a step-by-step process to tackle every possible problem faced while training deep models. You will implement convolutional and recurrent neural networks, adversarial networks, and more with the help of this handy guide. In addition to this, you will learn how to train these models for real-world image and language processing tasks. By the end of this book, you will have a practical, hands-on understanding of how you can leverage the power of Python and Keras to perform effective deep learning What you will learn Install and configure Keras in TensorFlow Master neural network programming using the Keras library Understand the different Keras layers Use Keras to implement simple feed-forward neural networks, CNNs and RNNs Work with various datasets and models used for image and text classification Develop text summarization and reinforcement learning models using Keras Who this book is for Keras Deep Learning Cookbook is for you if you are a data scientist or machine learning expert who wants to find practical solutions to common problems encountered while training deep learning models. A basic understanding of Python and some experience in machine learning and neural networks is required for this book.

This IBM® Redbooks® publication is a guide about the IBM PowerAI Deep Learning solution. This book provides an introduction to artificial intelligence (AI) and deep learning (DL), IBM PowerAI, and components of IBM PowerAI, deploying IBM PowerAI, guidelines for working with data and creating models, an introduction to IBM Spectrum™ Conductor Deep Learning Impact (DLI), and case scenarios. IBM PowerAI started as a package of software distributions of many of the major DL software frameworks for model training, such as TensorFlow, Caffe, Torch, Theano, and the associated libraries, such as CUDA Deep Neural Network (cuDNN). The IBM PowerAI software is optimized for performance by using the IBM Power Systems™ servers that are integrated with NVLink. The AI stack foundation starts with servers and accelerators. Graphical processing unit (GPU) accelerators are well-suited for the compute-intensive nature of DL training, and servers with the highest CPU to GPU bandwidth, such as IBM Power Systems servers, enable the high-performance data transfer that is required for larger and more complex DL models. This publication targets technical readers, including developers, IT specialists, systems architects, brand specialist, sales team, and anyone looking for a guide about how to understand the IBM PowerAI Deep Learning architecture, framework configuration, application and workload configuration, and user infrastructure.

This open access work presents selected results from the European research and innovation project IMPROVE which yielded novel data-based solutions to enhance machine reliability and efficiency in the fields of simulation and optimization, condition monitoring, alarm management, and quality prediction.

Practical LXC and LXD

Cara mudah desain sistem operasi Linux Ubuntu, 16.04 LTS edition dalam 5 jam

Pro SQL Server on Linux

Proceedings of the 12th International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS-2018)

Innovative Mobile and Internet Services in Ubiquitous Computing

Intelligent Methods for the Factory of the Future

Network and System Security

***Build your own hybrid cloud strategy with this comprehensive learning guide. Key Features Build a hybrid cloud strategy for your organization with AWS and OpenStack Leverage Hybrid Cloud to design a complex deployment pipeline Learn to implement security and monitoring best practices with real-world examples Book Description Hybrid cloud is currently the buzz word in the cloud world. Organizations are planning to adopt hybrid cloud strategy due to its advantages such as untested workloads, cloud-bursting, cloud service brokering and so on. This book will help you understand the dynamics, design principles, and deployment strategies of a Hybrid Cloud. You will start by understanding the concepts of hybrid cloud and the problems it solves as compared to a stand-alone public and private cloud. You will be delving into the different architecture and design of hybrid cloud. The book will then cover advanced concepts such as building a deployment pipeline, containerization strategy, and data storage mechanism. Next up, you will be able to deploy an external CMP to run a Hybrid cloud and integrate it with your OpenStack and AWS environments. You will also understand the strategy for designing a Hybrid Cloud using containerization and work with pre-built solutions like vCloud Air, VMware for AWS, and Azure Stack. Finally, the book will cover security and***

**monitoring related best practices that will help you secure your cloud infrastructure. By the end of the book, you will be in a position to build a hybrid cloud strategy for your organization. What you will learn Learn the demographics and definitions of Hybrid Cloud Understand the different architecture and design of Hybrid Cloud Explore multi-cloud strategy and use it with your hybrid cloud Implement a Hybrid Cloud using CMP / Common API's Implement a Hybrid Cloud using Containers Overcome various challenges and issues while working with your Hybrid Cloud Understand how to monitor your Hybrid Cloud Discover the security implications in the Hybrid Cloud Who this book is for This book is targeted at cloud architects, cloud solution providers, DevOps engineers, or any working stakeholder who wants to learn about the hybrid cloud architecture. A basic understanding of public and private cloud is desirable.**

**The two-volume set CCIS 827 and 828 constitutes the thoroughly refereed proceedings of the Third International Conference on Next Generation Computing Technologies, NGCT 2017, held in Dehradun, India, in October 2017. The 135 full papers presented were carefully reviewed and selected from 948 submissions. There were organized in topical sections named: Smart and Innovative Trends in Communication Protocols and Standards; Smart and Innovative Trends in Computational Intelligence and Data Science; Smart and Innovative Trends in Image Processing and Machine Vision; Smart Innovative Trends in Natural Language Processing for Indian Languages; Smart Innovative Trends in Security and Privacy.**

**Buku ini membahas mengenai cara kustomisasi sistem operasi linux, khususnya Ubuntu 16.04 LTS, dengan kustomisasi sistem operasi ini dapat digunakan untuk melakukan konfigurasi secara cepat sesuai kebutuhan pada saat melakukan instalasi sistem operasi.**

**The book covers the world's leading open source application platform: Red Hat Enterprise Linux (RHEL) 6 is a Linux distribution produced by Red Hat that was targeted toward the commercial market, including mainframes. Red Hat Enterprise Linux is released in server versions for x86, x86\_64, Itanium, PowerPC and IBM System z, and desktop versions for x86 and x86\_64. All of Red Hat's official support and training, and the Red Hat Certification Program center on the Red Hat Enterprise Linux platform. It next Covers CentOS 6 Server: is a community-supported, free and open source operating system based on Red Hat Enterprise Linux. It exists to provide a free enterprise class computing platform and strives to maintain 100% binary compatibility with its upstream distribution .CentOS stands for "Community ENTerprise Operating System". CentOS is the perfect server for people who need an enterprise class operating system stability without the cost of certification and support and pocket burning baggage that comes with proprietary software. And the beauty is CentOS is free. Finally, we end with Ubuntu 14.04 LTS Server which transformed the IT environments worldwide. Realize the full potential of your infrastructure with a reliable, easy-to-integrate technology platform. It's the leading platform for scale-out computing, Ubuntu Server helps you make the most of your infrastructure. Whether you want to deploy an OpenStack cloud, a Hadoop cluster or a 50,000-node render farm, Ubuntu Server delivers the best value scale-out performance available. It's an**

***excellent book for those who still use or have fond memory of these three industry standard distros.***

***Mastering Ubuntu Server***

***IMPROVE - Innovative Modelling Approaches for Production Systems to Raise Validatable Efficiency***

***Mastering Linux Servers***

***Google Cloud Platform Cookbook***

***200+ Automation Examples For Linux and Windows System Administrator and DevOps***

***Including Container-Based Deployment with Docker and Kubernetes***

***Start your Journey into Free and Open Source Software***

This IBM® Redbooks® publication describes the new member of the IBM Z® family, IBM z14™. IBM z14 is the trusted enterprise platform for pervasive encryption, integrating data, transactions, and insights into the data. A data-centric infrastructure must always be available with a 99.999% or better availability, have flawless data integrity, and be secured from misuse. It also must be an integrated infrastructure that can support new applications. Finally, it must have integrated capabilities that can provide new mobile capabilities with real-time analytics that are delivered by a secure cloud infrastructure. IBM z14 servers are designed with improved scalability, performance, security, resiliency, availability, and virtualization. The superscalar design allows z14 servers to deliver a record level of capacity over the prior IBM Z platforms. In its maximum configuration, z14 is powered by up to 170 client characterizable microprocessors (cores) running at 5.2 GHz. This configuration can run more than 146,000 million instructions per second (MIPS) and up to 32 TB of client memory. The IBM z14 Model M05 is estimated to provide up to 35% more total system capacity than the IBM z13® Model NE1. This Redbooks publication provides information about IBM z14 and its functions, features, and associated software support. More information is offered in areas that are relevant to technical planning. It is intended for systems engineers, consultants, planners, and anyone who wants to understand the IBM Z servers functions and plan for their usage. It is intended as an introduction to mainframes. Readers are expected to be generally familiar with existing IBM Z technology and terminology.

This collection of essays consists of selected papers presented at the 24th IEEE FRUCT conference. It highlights the most pressing research topics in infocommunication technologies, such as challenges in the development of next generation networks, the architectures and design of innovative knowledge-based systems, and innovations in healthcare and eHealth.

Training Manual: Covering Ubuntu 64.04 LTS Server Installation & Administration, Media servers: Plex, KODI, Emby & Madsonic; NAS Server: s FreeNAS & OpenMediaVault Media Storage

Training Manual: Covering Ubuntu 64.04 LTS Server, Debian 8, Application Servers: Apache Tomcat 8, JBoss-eap 6, GlassFish 4, Eclipse IDE, and Postfix with Dovecot, OpenLDAP, phpLDAPadmin

Pervasive Systems, Algorithms and Networks

Proceedings of the 13th International Conference on P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC-2018)

Ubuntu 16.04 Lts

Linux Containers for Virtualization and Orchestration

Emerging Topics and Questions in Infocommunication Technologies

Introducing Disaster Recovery with Microsoft Azure

Beginning Elastic Stack

**Ansible is a popular open-source IT automation technology for scripting applications in a wide variety of domains. It is free, portable, powerful, and remarkably easy and fun to use. This book is a tool to learn the Ansible automation technology with some real-life examples. Whenever you are new to automation or a professional automation engineer, this book's goal is to bring you quickly up to speed on the fundamentals of the core Ansible language. Every successful IT department needs automation nowadays for bare metal servers, virtual machines, cloud, containers, and edge computing. Automate your IT journey with Ansible automation technology. I'm going to teach you example by example how to accomplish the most common System Administrator tasks. You are going to start with the installation of Ansible in Windows 10 and Windows 11 and use the most command package manager and archives. Each of the 50+ lessons summarizes a module: from the most important parameter to some live demo of code and real-life usage. Each code is battle proved in the real life. Console interaction and verification are included in every video. A mundane activity like installing software, verifying a system is up-to-date, rebooting a server, installing Google Chrome, copying files from the local controller to a remote system, could be automated with some lines of code and these are only some of the long lists included in the course. There are some Ansible codes usable in all the Windows systems and some specific for Windows Server. The Ansible troubleshooting lessons teach you how to read the error message, how to reproduce, and the process of troubleshooting and resolution. Are you ready to automate your day with Ansible?**

**Get SQL Server up and running on the Linux operating system and containers. No database professional managing or developing SQL Server on Linux will want to be without this deep and authoritative guide by one of the most respected experts on SQL Server in the industry. Get an inside look at how SQL Server for Linux works through the eyes of an engineer on the team that made it possible. Microsoft SQL Server is one of the leading database platforms in the industry, and SQL Server 2017 offers developers and administrators the ability to run a database management system on Linux, offering proven support for enterprise-level features and without onerous licensing terms. Organizations invested in Microsoft and open source technologies are now able to run a unified database platform across all their operating system investments. Organizations are further able to take full advantage of containerization through popular platforms such as Docker and Kubernetes. Pro SQL Server on Linux walks you through installing and configuring SQL Server on the Linux platform. The author is one of the principal architects of SQL Server for Linux, and brings a corresponding depth of knowledge that no database professional or developer on Linux will want to be without. Throughout this book**



are internals of how SQL Server on Linux works including an in depth look at the innovative architecture. The book covers day-to-day management and troubleshooting, including diagnostics and monitoring, the use of containers to manage deployments, and the use of self-tuning and the in-memory capabilities. Also covered are performance capabilities, high availability, and disaster recovery along with security and encryption. The book covers the product-specific knowledge to bring SQL Server and its powerful features to life on the Linux platform, including coverage of containerization through Docker and Kubernetes. What You'll Learn Learn about the history and internal of the unique SQL Server on Linux architecture. Install and configure Microsoft's flagship database product on the Linux platform Manage your deployments using container technology through Docker and Kubernetes Know the basics of building databases, the T-SQL language, and developing applications against SQL Server on Linux Use tools and features to diagnose, manage, and monitor SQL Server on Linux Scale your application by learning the performance capabilities of SQL Server Deliver high availability and disaster recovery to ensure business continuity Secure your database from attack, and protect sensitive data through encryption Take advantage of powerful features such as Failover Clusters, Availability Groups, In-Memory Support, and SQL Server's Self-Tuning Engine Learn how to migrate your database from older releases of SQL Server and other database platforms such as Oracle and PostgreSQL Build and maintain schemas, and perform management tasks from both GUI and command line Who This Book Is For Developers and IT professionals who are new to SQL Server and wish to configure it on the Linux operating system. This book is also useful to those familiar with SQL Server on Windows who want to learn the unique aspects of managing SQL Server on the Linux platform and Docker containers. Readers should have a grasp of relational database concepts and be comfortable with the SQL language.

This book presents the latest research findings, methods and development techniques related to Ubiquitous and Pervasive Computing (UPC) as well as challenges and solutions from both theoretical and practical perspectives with an emphasis on innovative, mobile and internet services. With the proliferation of wireless technologies and electronic devices, there is a rapidly growing interest in Ubiquitous and Pervasive Computing (UPC). UPC makes it possible to create a human-oriented computing environment where computer chips are embedded in everyday objects and interact with physical world. It also allows users to be online even while moving around, providing them with almost permanent access to their preferred services. Along with a great potential to revolutionize our lives, UPC also poses new research challenges.

This book constitutes the refereed proceedings of the 16th International Symposium on Pervasive Systems, Algorithms and Networks, I-SPAN 2019, held in Naples, Italy, in September 2019. The 32 full papers and 8 short papers were carefully reviewed and selected from 89 submissions. The papers focus on all aspects of: big data analytics & machine learning; cyber security; cloud fog & edge computing; communication solutions; high performance computing and applications; consumer cyber security; and vehicular technology.

13th International Conference, NSS 2019, Sapporo, Japan, December 15–18, 2019, Proceedings

The Ultimate Linux Ubuntu 16.04 Powered Media Streaming & Storage Servers

Training Manual: Covering Application Servers: Apache Tomcat 9, Jboss-Eap 6, Glassfish 4, Eclipse Ide, and Backtrack 5 Pentest

## **Ubuntu 16.04 LTS Desktop: Applications and Administration**

### **IBM z14 ZR1 Technical Guide**

#### **Master The Art of Hyperledger Fabric on Kubernetes**

#### **Mastering Hyperledger Fabric**

This IBM® Redbooks® publication describes the new member of the IBM Z® family, IBM z14™ Model ZR1 (Machine Type 3907). It includes information about the Z environment and how it helps integrate data and transactions more securely, and can infuse insight for faster and more accurate business decisions. The z14 ZR1 is a state-of-the-art data and transaction system that delivers advanced capabilities, which are vital to any digital transformation. The z14 ZR1 is designed for enhanced modularity, in an industry standard footprint. A data-centric infrastructure must always be available with a 99.999% or better availability, have flawless data integrity, and be secured from misuse. It also must be an integrated infrastructure that can support new applications. Finally, it must have integrated capabilities that can provide new mobile capabilities with real-time analytics that are delivered by a secure cloud infrastructure. IBM z14 ZR1 servers are designed with improved scalability, performance, security, resiliency, availability, and virtualization. The superscalar design allows z14 ZR1 servers to deliver a record level of capacity over the previous IBM Z platforms. In its maximum configuration, z14 ZR1 is powered by up to 30 client characterizable microprocessors (cores) running at 4.5 GHz. This configuration can run more than 29,000 million instructions per second and up to 8 TB of client memory. The IBM z14 Model ZR1 is estimated to provide up to 54% more total system capacity than the IBM z13s® Model N20. This Redbooks publication provides information about IBM z14 ZR1 and its functions, features, and associated software support. More information is offered in areas that are relevant to technical planning. It is intended for systems engineers, consultants, planners, and anyone who wants to understand the IBM Z servers functions and plan for their usage. It is intended as an introduction to mainframes. Readers are expected to be generally familiar with IBM Z technology and terminology.

This IBM Redbooks publication describes how to implement an Open Platform for Database as a Service (DBaaS) on IBM Power Systems environment for Linux, and demonstrate the open source tools, optimization and best practices guidelines for it. Open Platform for DBaaS on Power Systems is an on-demand, secure, and scalable self-service database platform that automates provisioning and administration of databases to support new business applications and information insights. This publication addresses topics to help sellers, architects, brand specialists, distributors, resellers and anyone offering secure and scalable Open Platform for DBaaS on Power Systems solution with APIs that are consistent across heterogeneous open database types. An Open Platform for DBaaS on Power Systems solution has the capability to accelerate business success by providing an infrastructure, and tools leveraging Open Source and OpenStack software engineered to optimize hardware and software between workloads and resources so you have a responsive, and an adaptive environment. Moreover, this publication provides documentation to transfer the how-to-skills for cloud oriented operational management of Open Platform for DBaaS on Power Systems service and underlying infrastructure to the technical teams. Open Platform for DBaaS on Power Systems mission is to

provide scalable and reliable cloud database as a service provisioning functionality for both relational and non-relational database engines, and to continue to improve its fully-featured and extensible open source framework. For example, Trove is a database as a service for OpenStack. It is designed to run entirely on OpenStack, with the goal of allowing users to quickly and easily utilize the features of a relational or non-relational database without the burden of handling complex administrative tasks. Cloud users and database administrators can provision and manage multiple database instances as needed. Initially, the service focuses on providing resource isolation at high performance while automating complex administrative tasks including deployment, configuration, patching, backups, restores, and monitoring. In the context of this publication, the monitoring tool implemented is Nagios Core which is an open source monitoring tool. Hence, when you see a reference of Nagios in this book, Nagios Core is the open source monitoring solution implemented. Also note that the implementation of Open Platform for DBaaS on IBM Power Systems is based on open source solutions. This book is targeted toward sellers, architects, brand specialists, distributors, resellers and anyone developing and implementing Open Platform for DBaaS on Power Systems solutions.

Implement a SOHO or SMB Linux infrastructure to expand your business and associated IT capabilities. Backed by the expertise and experienced guidance of the authors, this book provides everything you need to move your business forward. Pro Linux System Administration makes it easy for small- to medium-sized businesses to enter the world of zero-cost software running on Linux and covers all the distros you might want to use, including Red Hat, Ubuntu, Debian, and CentOS. Pro Linux System Administration takes a layered, component-based approach to open source business systems, while training system administrators as the builders of business infrastructure. Completely updated for this second edition, Dennis Matotek takes you through an infrastructure-as-code approach, seamlessly taking you through steps along the journey of Linux administration with all you need to master complex systems. This edition now includes Jenkins, Ansible, Logstash and more. What You'll Learn: Understand Linux architecture Build, back up, and recover Linux servers Create basic networks and network services with Linux Build and implement Linux infrastructure and services including mail, web, databases, and file and print Implement Linux security Resolve Linux performance and capacity planning issues Who This Book Is For: Small to medium-sized business owners looking to run their own IT, system administrators considering migrating to Linux, and IT systems integrators looking for an extensible Linux infrastructure management approach.

Use Linux containers as an alternative virtualization technique to virtualize your operating system environment. This book will cover LXC ' s unmatched flexibility with virtualization and LXD ' s smooth user experience. Practical LXC and LXD begins by introducing you to Linux containers (LXC and LXD). You will then go through use cases based on LXC and LXD. Next, you will see the internal workings of LXC and LXD by considering the repositories and templates used. You will then learn how to integrate LXC and LXD with common virtualization and orchestration tools such as libvirt and SaltStack. Finally, you will dive into containerization and security. The book will explore some of the common problems in security and provide a case study on how containerization can help mitigate some of the operating system-level security issues in an IoT environment. What You Will Learn

Get an introduction to Linux containers Discover the basics of LXC and LXD See use cases that can be solved with LXC and LXD – for developers, devops, and system administrators Master LXC and LXD repositories Use LXC and LXD with common virtualization and orchestration tools Consider a containerization and security in IoT case study Who This Book Is For The audience for this book should have basic knowledge of Linux and software development in general. The intended readership is primarily software developers, operations engineers, and system administrators who are interested in devops, though managers and enthusiasts will also benefit from this book.

Google Cloud Platform Administration

IBM Open Platform for DBaaS on IBM Power Systems

Keras Deep Learning Cookbook

ICICT 2021, London, Volume 1

National Cyber Summit (NCS) Research Track

Ubuntu Server 16. 04 LTS Percise

Ansible For Windows By Examples

Mastering Hyperledger Fabric. A one-stop solution to become Master in the Hyperledger Fabric Key Features Detailed Explanation of One way TLS and mutual TLS Detailed Explanation of docker sockets (docker.sock) Exposed functionalities of Fabric CLI's and SDK's Enterprise-level chaincode development A glimpse of Hyperledger Fabric 2.0 Advanced examples of Node and golang Fabric SDK Onboard new organization using Node.js SDK (No more CLI) CI/CD for chaincode (Install chaincode directly onto peers from GitHub using Node.js) Fabric setup explanation with Different real-time use cases Deployment of Hyperledger Fabric using docker swarm and Kubernetes Setup and configure caliper to check benchmarks Monitor consortium with Prometheus and grafana Monitor docker and docker swarm using swarmpit and logspout Logging consortium with ELK/EFK stack Some interesting open-source tools and some Bonus concepts Table of Contents Chapter1: Introduction to the Hyperledger Landscape Chapter2: The Disruptive Potential of TLS Chapter3: All about docker sockets Chapter4: Installation Guide Of Prerequisites Chapter5: All about fabric CLI Chapter6: All about SDK's (go lang and Node.js) Chapter7: Advanced Chaincode Development Chapter8: End to End fabric consortium with Solo consensus using docker with one use case Chapter9: End to End fabric Consortium with Kafka consensus using docker swarm with one use case Chapter10: End to End fabric Consortium with Raft consensus using Kubernetes with one use case Chapter11: Private Data Concepts, Consortium level ACL(Access Control Lists) and raft consensus mechanism Chapter12: Setup and Benchmark Blockchain Consortium Using Caliper Chapter13: Monitoring Consortium with Prometheus and grafana Chapter14: Logging Consortium with ELK Stack Chapter15: Glimpse of Hyperledger fabric 2.0 Chapter16: Some Interesting tools Who this Book is For This Book benefits Software Engineers who are ready to shift their focus to distributed technologies and Blockchain. This book provides a comprehensive view of Solution Architecture, so it will be easy for architects to architect their solution. CTO's around the world want to add hyperledger fabric to their technology stack. Managers to

cope up with the latest trend. Faculty Professors in order to get industry insights. Even Engineering Students who want to be ready with the latest technologies. Book Description Mastering Hyperledger Fabric is a craving topic for all Hyperledger Fabric Developers around the world. Hyperledger Fabric is an open-source project that helps organizations create and maintain permissioned distributed Blockchain consortiums. This book is for readers who are looking for Hyperledger offerings to build end-to-end projects with growing complexity and functionalities. This book will be a one-stop solution for all developers who want to build blockchain consortiums using Hyperledger Fabric. Topics include TLS, Unix sockets, caliper(Benchmark tool), raft consensus, advanced chaincode development, key collision and MVCC, chaincode access controls, chaincode encryption, node.js SDK, goLang SDK, docker daemon API, private data concepts, onboarding organizations using node.js SDK, deploy hyperledger fabric using Kubernetes, deploy hyperledger fabric using docker swarm, monitoring hyperledger fabric, monitoring Kubernetes, monitoring docker swarm, logging hyperledger fabric. After reading this book the reader will be able to set up Production grade hyperledger fabric consortium using raft consensus mechanisms with monitoring using Prometheus and grafana, even logging. This book explains so many key concepts of hyperledger fabric including 2.0 and written with three years of hyperledger fabric production experience.

Ansible is an Open Source IT automation tool. This book contains all of the obvious and not-so-obvious best practices of Ansible automation for Security and Compliance. Every successful IT department needs automation nowadays for bare metal servers, virtual machines, cloud, containers, and edge computing. Automate your IT journey with Ansible automation technology. You are going to start with the installation of Ansible in Enterprise Linux, Community Linux, Windows, and macOS using the most command package manager and archives. Each of the 100+ lessons summarizes a module: from the most important parameter to some Ansible code and real-life usage. Each code is battle proved in the real life. Simplifying mundane activities like creating a text file, extracting and archiving, fetching a repository using HTTPS or SSH connections could be automated with some lines of code and these are only some of the long lists included in the course. There are some Ansible codes usable in all the Linux systems, some specific for RedHat-like, Debian-like, and Windows systems. The 20+ Ansible troubleshooting lesson teaches you how to read the error message, how to reproduce, and the process of troubleshooting and resolution. Are you ready to automate your day with Ansible? Examples in the book are tested with the latest version of Ansible 2.9+ and Ansible Core 2.11+.

This book constitutes the proceedings of the 13th International Conference on Network and System Security, NSS 2019, held in Sapporo, Japan, in December 2019. The 36 full papers and 7 short papers presented together with 4 invited papers in this book were carefully reviewed and selected from 89 initial submissions. The papers cover a wide range of topics in the field, including authentication, access control, availability, integrity, privacy, confidentiality, dependability and sustainability of computer networks and systems.

Get up to date with the finer points of Ubuntu Server using this comprehensive guide About This Book Get well-versed with newly-added features in Ubuntu 16.04 Master the art of installing, managing, and troubleshooting Ubuntu Server A practical easy-to-

understand book that will help you enhance your existing skills. Who This Book Is For This book is intended for readers with intermediate or advanced-beginner skills with Linux, who would like to learn all about setting up servers with Ubuntu Server. This book assumes that the reader knows the basics of Linux, such as editing configuration files and running basic commands. What You Will Learn Learn how to manage users, groups, and permissions Encrypt and decrypt disks with Linux Unified Key Setup /Luks Setup SSH for remote access, and connect it to other nodes Understand how to add, remove, and search for packages Use NFS and Samba to share directories with other users Get to know techniques for managing Apache and MariaDB Explore best practices and troubleshooting techniques In Detail Ubuntu is a Debian-based Linux operating system, and has various versions targeted at servers, desktops, phones, tablets and televisions. The Ubuntu Server Edition, also called Ubuntu Server, offers support for several common configurations, and also simplifies common Linux server deployment processes. With this book as their guide, readers will be able to configure and deploy Ubuntu Servers using Ubuntu Server 16.04, with all the skills necessary to manage real servers. The book begins with the concept of user management, group management, as well as file-system permissions. To manage your storage on Ubuntu Server systems, you will learn how to add and format storage and view disk usage. Later, you will also learn how to configure network interfaces, manage IP addresses, deploy Network Manager in order to connect to networks, and manage network interfaces. Furthermore, you will understand how to start and stop services so that you can manage running processes on Linux servers. The book will then demonstrate how to access and share files to or from Ubuntu Servers. You will learn how to create and manage databases using MariaDB and share web content with Apache. To virtualize hosts and applications, you will be shown how to set up KVM/Qemu and Docker and manage virtual machines with virt-manager. Lastly, you will explore best practices and troubleshooting techniques when working with Ubuntu Servers. By the end of the book, you will be an expert Ubuntu Server user well-versed in its advanced concepts. Style and Approach This book is an advanced guide that will show readers how to administer, manage, and deploy Ubuntu server and will also provide expert-level knowledge on advanced security and backup techniques.

Quick Start Guide to Ubuntu Server

Hybrid Cloud for Architects

Beginning Ubuntu for Windows and Mac Users

Pro Linux System Administration

The Virtualization Cookbook for IBM z Systems Volume 4: Ubuntu Server 16.04

Advances on P2P, Parallel, Grid, Cloud and Internet Computing

Windows Server 2016 Inside Out (includes Current Book Service)

*Ubuntu Server 16. 04 LTS PerciseAdministration and Quick Reference*

*Learn how to install, configure and implement the Elastic Stack (Elasticsearch, Logstash and Kibana) - the invaluable tool for anyone deploying a centralized log management solution for*

servers and apps. You will see how to use and configure Elastic Stack independently and alongside Puppet. Each chapter includes real-world examples and practical troubleshooting tips, enabling you to get up and running with Elastic Stack in record time. Fully customizable and easy to use, Elastic Stack enables you to be on top of your servers all the time, and resolve problems for your clients as fast as possible. Supported by Puppet and available with various plugins. Get started with *Beginning Elastic Stack* today and see why many consider Elastic Stack the best option for server log management. What You Will Learn: Install and configure Logstash Use Logstash with Elasticsearch and Kibana Use Logstash with Puppet and Foreman Centralize data processing Who This Book Is For: Anyone working on multiple servers who needs to search their logs using a web interface. It is ideal for server administrators who have just started their job and need to look after multiple servers efficiently.

Practical recipes to implement cost-effective and scalable cloud solutions for your organization Key Features Implement Google Cloud services in your organization Leverage Google Cloud components to secure your organization's data A recipe-based guide that promises hands-on experience in deploying a highly scalable and available environment Book Description Google Cloud Platform is a cloud computing platform that offers products and services to host applications using state-of-the-art infrastructure and technology. You can build and host applications and websites, store data, and analyze data on Google's scalable infrastructure. This book follows a recipe-based approach, giving you hands-on experience to make the most of Google Cloud services. This book starts with practical recipes that explain how to utilize Google Cloud's common services. Then, you'll see how to make full use of Google Cloud components such as networking, security, management, and developer tools. Next, we'll deep dive into implementing core Google Cloud services into your organization, with practical recipes on App Engine, Compute Engine microservices with Cloud Functions, virtual networks, and Cloud Storage. Later, we'll provide recipes on implementing authentication and security, Cloud APIs, command-line management, deployment management, and the Cloud SDK. Finally, we'll cover administration troubleshooting tasks with the Compute and Container Engines and we'll show how to monitor your organization's efficiency with best practices. By the end of this book, you'll have a complete understanding of how to implement Google Cloud services in your organization with ease. What you will learn Host a Python application on Google Compute Engine Host an application using Google Cloud Functions Migrate a MySQL DB to Cloud Spanner Configure a network for a highly available

*application on GCP Learn simple image processing using Storage and Cloud Functions Automate security checks using Policy Scanner Understand tools for monitoring a production environment in GCP Learn to manage multiple projects using service accounts Who this book is for This book is for IT professionals, engineers, and developers looking at implementing Google Cloud in their organizations. Administrators and architects planning to make their organization more efficient with Google Cloud will also find this book useful. Basic understanding of Cloud services and the Google Cloud platform is necessary.*

*Discover how to get the most out of Ubuntu for work, home, and play. Learning a new operating system can feel daunting, especially if you're used to Windows or OS X. If you've been afraid to try Ubuntu because you don't know where to start, this book introduces you to a wide selection of software and settings that will make your computer ready to work for you. You'll see how Ubuntu can make your computing life easy. In addition to a tour of Ubuntu's modern and easy-to-use interface, you'll also learn how Ubuntu's Software Updater keeps all of your software secure and up-to-date. Browsing the Internet becomes faster and safer. Creating documents and sharing with others is built right in. Enjoying your music and movie libraries helps you unwind. Ubuntu is the world's third most popular operating system and powers desktop and laptop computers, servers, private and public clouds, and embedded devices. There's never been a better time to install Ubuntu and move to an open source way of life. Completely updated for this exciting second edition, Beginning Ubuntu for Windows and Mac Users will help you start your journey into Free and Open Source Software with Ubuntu 16.04 LTS. What You'll Learn Understand the advantages of Ubuntu and its variants—Kubuntu, Xubuntu, and more Install Ubuntu on its own or alongside your computer's existing operating system Search Ubuntu's catalog of thousands of applications—all ready to install with a single click Work with files and disks that were created with Windows and OS X Run simple, interesting tasks and games using the command line Customize Ubuntu in powerful ways and get work done with virtual machines Who This Book Is For Anyone who wants to move to using an open source operating system.*

*Expert Apache Cassandra Administration*

*Build robust hybrid cloud solutions using AWS and OpenStack*

*50+ Automation Examples For Windows System Administrator And DevOps*

*Implement, deploy, maintain, and migrate applications on Google Cloud Platform*

*Ansible For Security by Examples*



*Design highly available, scalable, and secure cloud solutions on GCP*

*31st British International Conference on Databases, BICOD 2017, London, UK, July 10-12, 2017, Proceedings*

***Conquer Windows Server 2016—from the inside out! Dive into Windows Server 2016—and really put your Windows Server expertise to work. Focusing on Windows Server 2016’s most powerful and innovative features, this supremely organized reference packs hundreds of timesaving solutions, tips, and workarounds—all you need to plan, implement, or manage Windows Server in enterprise, data center, cloud, and hybrid environments. Fully reflecting Windows Server new capabilities for the cloud-first era, Orin covers everything from Nano Server to Windows Server and Hyper-V Containers. You’ll discover how experts tackle today’s essential tasks—and challenge yourself to new levels of mastery. • Optimize the full Windows Server 2016 lifecycle, from planning and configuration through rollout and administration • Ensure fast, reliable upgrades and migrations • Seamlessly deliver core DNS, DHCP, file, print, storage, and Internet services • Use IPAM to centrally manage all enterprise DNS and DHCP infrastructure • Gain dramatic storage utilization improvements with built-in deduplication and storage replica • Build flexible cloud and hybrid environments with Windows Containers and Shielded VMs • Seamlessly integrate Azure IaaS services with Windows Server 2016 • Slash resource usage and improve availability with tiny Nano Server installations • Improve configuration management with Desired State Configuration and Chef • Deliver Active Directory identity, certificate, federation, and rights management services • Protect servers, clients, assets, and users with advanced Windows Server 2016 security features including Just Enough Administration For Experienced Windows Server Users and IT Professionals • Your role: Experienced intermediate-to-advanced level Windows Server user or IT professional • Prerequisites: Basic understanding of Windows Server procedures, techniques, and navigation***

***This book presents the latest research findings, as well as innovative theoretical and practical research results, methods and development techniques related to P2P, grid, cloud and Internet computing. It also reveals the synergies among such large scale computing paradigms. P2P, Grid, Cloud and Internet computing technologies have rapidly become established as breakthrough paradigms for solving complex problems by enabling aggregation and sharing of an increasing variety of distributed computational resources on a large scale. Grid computing originated as a paradigm for high-performance computing, offering an alternative to expensive supercomputers through different forms of large-scale distributed computing. P2P computing emerged as a new paradigm following on from client-server and web-based computing and has proved useful in the development of social networking, B2B (Business to Business), B2C (Business to Consumer), B2G (Business to Government), and B2E (Business to Employee). Cloud computing has been described as a “computing paradigm where the boundaries of computing are determined by economic rationale rather than technical limits”. Cloud computing has fast become the computing paradigm with applicability and adoption in all***

**domains and providing utility computing at large scale. Lastly, Internet computing is the basis of any large-scale distributed computing paradigm; it has very quickly developed into a vast and flourishing field with enormous impact on today's information societies and serving as a universal platform comprising a large variety of computing forms such as grid, P2P, cloud and mobile computing.**

**Ansible is an Open Source IT automation tool. This book contains all of the obvious and not-so-obvious best practices of Ansible automation. Every successful IT department needs automation nowadays for bare metal servers, virtual machines, cloud, containers, and edge computing. Automate your IT journey with Ansible automation technology. You are going to start with the installation of Ansible in Enterprise Linux, Community Linux, Windows, and macOS using the most command package manager and archives. Each of the 200+ lessons summarizes a module: from the most important parameter to some Ansible code and real-life usage. Each code is battle proved in the real life. Simplifying mundane activities like creating a text file, extracting and archiving, fetching a repository using HTTPS or SSH connections could be automated with some lines of code and these are only some of the long lists included in the course. There are some Ansible codes usable in all the Linux systems, some specific for RedHat-like, Debian-like, and Windows systems. The 20+ Ansible troubleshooting lesson teaches you how to read the error message, how to reproduce, and the process of troubleshooting and resolution. Are you ready to automate your day with Ansible? Examples in the book are tested with the latest version of Ansible 2.9+ and Ansible Core 2.11+.**

**Make the most of GCP's offerings to manage your data center workload and optimize deployments Key Features Discover new techniques to administer, manage, and deploy applications on GCP Understand effective solutions for storing, retrieving, and deploying your container images Explore various offerings of GCP for operations and security Book Description On-premise data centers are costly to manage. If you need a data center but don't want to deal with a physical one, Google Cloud Platform (GCP) is the solution. With GCP, you can build, test, and deploy applications on Google's infrastructure. Google Cloud Platform Administration begins with GCP fundamentals, with the help of which you will deploy your first app and gain an understanding of Google Cloud architecture and services. Furthermore, you will learn how to manage Compute, networking, and storage resources. As you make your way through the book, you will learn how to track and manage GCP's usage, monitoring, and billing access control. You will also be able to manage your GCP's access and permissions. In the concluding chapters, you will explore a list of different developer tools for managing and interacting with the GCP platform. By the end of this book, you will have learned how to effectively deploy workloads on GCP. What you will learn Understand all GCP Compute components Deploy and manage multiple GCP storage options Manage and utilize the networking resources offered by GCP Explore the functionalities and features of the GCP Container Understand the workings of GCP operations such as monitoring and error reporting Discover an immune GCP using its identity and security options Who this book is for Google Cloud Platform Administration is for**

*administrators, cloud architects, and engineers who want to leverage the upcoming Google Cloud Platform. Some basic understanding of cloud computing will be useful.*

*IBM z14 (3906) Technical Guide*

*Training Manual:RHLE7, CentOS-7, Ubuntu 14. 04LTS, Also Covering Fedora 23 Server, Mandriva Enterprise Server, OpenSUSE 13. 2, Ubuntu 16. 04 LTS, Debian 13*

*IBM PowerAI: Deep Learning Unleashed on IBM Power Systems Servers*