

Red Hat Ceph Storage

Eventually, you will utterly discover a further experience and ability by spending more cash. still when? reach you bow to that you require to get those all needs afterward having significantly cash? Why dont you try to acquire something basic in the beginning? Thats something that will lead you to understand even more concerning the globe, experience, some places, afterward history, amusement, and a lot more?

It is your definitely own become old to proceed reviewing habit. in the middle of guides you could enjoy now is **Red Hat Ceph Storage** below.

[OpenStack Operations Guide](#) - Tom Fifield
2014-04-24

Design, deploy, and maintain your own private or public Infrastructure as a Service (IaaS), using the open source OpenStack platform. In this practical guide, experienced developers and OpenStack contributors show you how to build clouds based on reference architectures, as well as how to perform daily administration tasks. Designed for horizontal scalability, OpenStack lets you build a cloud by integrating several technologies. This approach provides flexibility, but knowing which options to use can be bewildering. Once you complete this book, you'll know the right questions to ask while you organize compute, storage, and networking resources. If you already know how to manage multiple Ubuntu machines and maintain MySQL, you're ready to: Set up automated deployment and configuration Design a single-node cloud controller Use metrics to improve scalability Explore compute nodes, network design, and storage Install OpenStack packages Use an example architecture to help simplify decision-making Build a working environment to explore an IaaS cloud Manage users, projects, and quotas Tackle maintenance, debugging, and network troubleshooting Monitor, log, backup, and restore

[Software Defined Data Center with Red Hat Cloud and Open Source IT Operations Management](#) - Dino Quintero 2020-11-04

This IBM® Redbooks® publication delivers a Site Reliability Engineering (SRE) solution for cloud workloads that uses Red Hat OpenStack for Infrastructure as a Service (IaaS), Red Hat OpenShift for Platform as a Service (PaaS), and

IT operations management that uses open source tools. Today, customers are no longer living in a world of licensed software. Curiosity increased the demand for investigating the Open Source world for Community Open Source and Enterprise grade applications. IBM as one of the contributors to the Open Source community is interested in helping the software be maintained and supported. Having companies, such as IBM, support the evolution of Open Source software helps to keep the Open Source community striving for enterprise grade open source solutions. Lately, companies are working on deciphering how to take advantage of Enterprise and Community Open Source to implement in their enterprises. The business case for open source software is no longer a mystery and no surprise that most of the new positions in IT enterprises are related to open source projects. The ability of a large enterprise to manage this sort of implementations is to engage in a hypertrophied cooperation, where the ability to not only cooperate with teams and people outside your organization, but also to find new ways of working together and devise new ways to improve the software and its code. A goal for this publication is to help the client's journey into the open source space and implement a private Cloud Container-based architecture with the ability to manage the entire IT Service Management processes from the open source framework. This publication describes the architecture and implementation details of the solution. Although not every piece of this solution is documented here, this book does provide instructions for what was achieved incorporating open source technologies.

Moreover, with this publication, the team shares their collaboration experiences working in a team of technologists, open source developers, Red Hat, and the open source community. This publication is for designers, developers, managers, and anyone who is considering starting a Cloud open source project, or users who started that journey. This book also can be a manual to guide the implementation of a technical viable architecture and help those enterprises participate in an open source project but have not done so before. The reader must be familiar with principles in programming and basic software engineering concepts, such as source code, compilers, and patches.

Learning Ceph - Karan Singh 2015-01-30

If you already have basic knowledge of GNU/Linux and storage systems, but have no experience of software-defined storage solutions and Ceph, and are eager to learn about it, this is the book for you. If you are looking for your next career jump as a Ceph administrator, this book is also ideal for you.

Trino: The Definitive Guide - Matt Fuller 2021-04-14

Perform fast interactive analytics against different data sources using the Trino high-performance distributed SQL query engine. With this practical guide, you'll learn how to conduct analytics on data where it lives, whether it's Hive, Cassandra, a relational database, or a proprietary data store. Analysts, software engineers, and production engineers will learn how to manage, use, and even develop with Trino. Initially developed by Facebook, open source Trino is now used by Netflix, Airbnb, LinkedIn, Twitter, Uber, and many other companies. Matt Fuller, Manfred Moser, and Martin Traverso show you how a single Trino query can combine data from multiple sources to allow for analytics across your entire organization. Get started: Explore Trino's use cases and learn about tools that will help you connect to Trino and query data Go deeper: Learn Trino's internal workings, including how to connect to and query data sources with support for SQL statements, operators, functions, and more Put Trino in production: Secure Trino, monitor workloads, tune queries, and connect more applications; learn how other organizations apply Trino

AWS System Administration - Mike Ryan 2018-08-08

With platforms designed for rapid adaptation and failure recovery such as Amazon Web Services, cloud computing is more like programming than traditional system administration. Tools for automatic scaling and instance replacement allow even small DevOps teams to manage massively scalable application infrastructures—if team members drop their old views of development and operations and start mastering automation. This comprehensive guide shows developers and system administrators how to configure and manage AWS services including EC2, CloudFormation, Elastic Load Balancing, S3, and Route 53. Sysadmins will learn will learn to automate their favorite tools and processes; developers will pick up enough ops knowledge to build a robust and resilient AWS application infrastructure. Launch instances with EC2 or CloudFormation Securely deploy and manage your applications with AWS tools Learn to automate AWS configuration management with Python and Puppet Deploy applications with Auto Scaling and Elastic Load Balancing Explore approaches for deploying application and infrastructure updates Save time on development and operations with reusable components Learn strategies for managing log files in AWS environments Configure a cloud-aware DNS service with Route 53 Use AWS CloudWatch to monitor your infrastructure and applications

Understanding and Deploying LDAP

Directory Services - Tim Howes 2003

Lightweight Directory Access Protocol (LDAP) is the standard for directory information access and is the underlying protocol for a variety of email systems, Web systems, and enterprise applications. LDAP enables central management of users, groups, devices, and other data, thereby simplifying directory management and reducing the total cost of ownership. Understanding and Deploying LDAP Directory Services, written by the creators of the protocol, is known as the LDAP bible and is the classic text for learning about LDAP and how to utilize it effectively. The Second Edition builds on this success by acting as an exhaustive resource for designing, deploying, and maintaining LDAP directory services. Topics such as

implementation pitfalls, establishing and maintaining user access to information, troubleshooting, and real-world scenarios will be thoroughly explored.

Ceph Cookbook - Vikhyat Umrao 2017-11-24

Over 100 effective recipes to help you design, implement, and troubleshoot manage the software-defined and massively scalable Ceph storage system. About This Book Implement a Ceph cluster successfully and learn to manage it. Recipe based approach in learning the most efficient software defined storage system Implement best practices on improving efficiency and security of your storage cluster Learn to troubleshoot common issues experienced in a Ceph cluster Who This Book Is For This book is targeted at storage and cloud engineers, system administrators, or anyone who is interested in building software defined storage, to power your cloud or virtual infrastructure. If you have basic knowledge of GNU/Linux and storage systems, with no experience of software defined storage solutions and Ceph, but eager to learn then this book is for you What You Will Learn Understand, install, configure, and manage the Ceph storage system Get to grips with performance tuning and benchmarking, and learn practical tips to help run Ceph in production Integrate Ceph with OpenStack Cinder, Glance, and Nova components Deep dive into Ceph object storage, including S3, Swift, and Keystone integration Configure a disaster recovery solution with a Ceph Multi-Site V2 gateway setup and RADOS Block Device mirroring Gain hands-on experience with Ceph Metrics and VSM for cluster monitoring Familiarize yourself with Ceph operations such as maintenance, monitoring, and troubleshooting Understand advanced topics including erasure-coding, CRUSH map, cache pool, and general Ceph cluster maintenance In Detail Ceph is a unified distributed storage system designed for reliability and scalability. This technology has been transforming the software-defined storage industry and is evolving rapidly as a leader with its wide range of support for popular cloud platforms such as OpenStack, and CloudStack, and also for virtualized platforms. Ceph is backed by Red Hat and has been developed by community of developers which has gained

immense traction in recent years. This book will guide you right from the basics of Ceph , such as creating blocks, object storage, and filesystem access, to advanced concepts such as cloud integration solutions. The book will also cover practical and easy to implement recipes on CephFS, RGW, and RBD with respect to the major stable release of Ceph Jewel. Towards the end of the book, recipes based on troubleshooting and best practices will help you get to grips with managing Ceph storage in a production environment. By the end of this book, you will have practical, hands-on experience of using Ceph efficiently for your storage requirements. Style and approach This step-by-step guide is filled with practical tutorials, making complex scenarios easy to understand.

DevOps with OpenShift - Stefano Picozzi
2017-07-10

For many organizations, a big part of DevOps' appeal is software automation using infrastructure-as-code techniques. This book presents developers, architects, and infra-ops engineers with a more practical option. You'll learn how a container-centric approach from OpenShift, Red Hat's cloud-based PaaS, can help your team deliver quality software through a self-service view of IT infrastructure. Three OpenShift experts at Red Hat explain how to configure Docker application containers and the Kubernetes cluster manager with OpenShift's developer- and operational-centric tools. Discover how this infrastructure-agnostic container management platform can help companies navigate the murky area where infrastructure-as-code ends and application automation begins. Get an application-centric view of automation—and understand why it's important Learn patterns and practical examples for managing continuous deployments such as rolling, A/B, blue-green, and canary Implement continuous integration pipelines with OpenShift's Jenkins capability Explore mechanisms for separating and managing configuration from static runtime software Learn how to use and customize OpenShift's source-to-image capability Delve into management and operational considerations when working with OpenShift-based application workloads Install a self-contained local version of the OpenShift environment on your computer

OpenStack for Architects - Michael Solberg
2017-02-06

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.

Advanced Platform Development with Kubernetes - Craig Johnston 2020-10-05

Leverage Kubernetes for the rapid adoption of emerging technologies. Kubernetes is the future of enterprise platform development and has become the most popular, and often considered the most robust, container orchestration system available today. This book focuses on platforming technologies that power the Internet of Things, Blockchain, Machine Learning, and the many layers of data and application management supporting them. Advanced Platform Development with Kubernetes takes you through the process of building platforms with these in-demand capabilities. You'll progress through the development of Serverless, CI/CD integration, data processing pipelines, event queues, distributed query engines, modern data warehouses, data lakes, distributed object storage, indexing and analytics, data routing and transformation, query engines, and data science/machine learning environments. You'll also see how to implement and tie together numerous essential and trending technologies including: Kafka, NiFi, Airflow, Hive, Keycloak, Cassandra, MySQL, Zookeeper, Mosquitto, Elasticsearch, Logstash, Kibana, Presto, Mino, OpenFaaS, and Ethereum. The book uses Golang and Python to demonstrate the development integration of custom container and Serverless functions, including interaction with the Kubernetes API. The exercises throughout teach Kubernetes through the lens of platform development, expressing the power and flexibility of Kubernetes with clear and pragmatic examples. Discover why Kubernetes is an excellent choice for any individual or organization looking to embark on developing a successful data and application platform. What You'll Learn Configure and install Kubernetes and k3s on vendor-neutral platforms, including

generic virtual machines and bare metal
Implement an integrated development toolchain
for continuous integration and deployment Use
data pipelines with MQTT, NiFi, Logstash, Kafka
and Elasticsearch Install a serverless platform
with OpenFaaS Explore blockchain network
capabilities with Ethereum Support a multi-
tenant data science platform and web IDE with
JupyterHub, MLflow and Seldon Core Build a
hybrid cluster, securely bridging on-premise and
cloud-based Kubernetes nodes Who This Book Is
For System and software architects, full-stack
developers, programmers, and DevOps
engineers with some experience building and
using containers. This book also targets readers
who have started with Kubernetes and need to
progress from a basic understanding of the
technology and "Hello World" example to more
productive, career-building projects.

Mastering Ceph - Nick Fisk 2017-05-30

Deep dive into the unified, distributed storage
system in order to provide excellent
performance About This Book Leverage Ceph's
advanced features such as erasure coding,
tiering, and Bluestore Solve large-scale
problems with Ceph as a tool by understanding
its strengths and weaknesses to develop the best
solutions A practical guide that covers engaging
use cases to help you use advanced features of
Ceph effectively Who This Book Is For If you are
a developer and an administrator who has
deployed a Ceph cluster before and are curious
about some of the most advanced features in
order to improve performance then this book is
for you What You Will Learn Know when and
how to use some of Ceph's advanced new
features Set up a test cluster with Ansible and
some virtual machines using VirtualBox and
Vagrant Develop novel solutions to massive
problems with librados and shared object
classes. Choose intelligent parameters for an
erasure coded pool and set it up. Configure the
Bluestore settings and see how they interact
with different hardware configurations. Keep
Ceph running through thick and thin with
tuning, monitoring and disaster recovery advice.
In Detail Mastering Ceph covers all that you
need to know to use Ceph effectively. Starting
with design goals and planning steps that should
be undertaken to ensure successful
deployments, you will be guided through to

setting up and deploying the Ceph cluster, with
the help of orchestration tools. Key areas of
Ceph including Bluestore, Erasure coding and
cache tiering will be covered with help of
examples. Development of applications which
use Librados and Distributed computations with
shared object classes are also covered. A section
on tuning will take you through the process of
optimizing both Ceph and its supporting
infrastructure. Finally, you will learn to
troubleshoot issues and handle various scenarios
where Ceph is likely not to recover on its own.
By the end of the book, you will be able to
successfully deploy and operate a resilient high
performance Ceph cluster. Style and Approach A
practical guide which has each chapter
explaining the concept, sharing tips and tricks
and a use case to implement the most powerful
features of Ceph

Mastering CloudForms Automation - Peter
McGowan 2016-08-22

Learn how to work with the Automate feature of
CloudForms, the powerful Red Hat cloud
management platform that lets you administer
your virtual infrastructure, including hybrid
public and private clouds. This practical hands-
on introduction shows you how to increase your
operational efficiency by automating day-to-day
tasks that now require manual input.
Throughout the book, author Peter McGowan
provides a combination of theoretical
information and practical coding examples to
help you learn the Automate object model. With
this CloudForms feature, you can create auto-
scalable cloud applications, eliminate manual
decisions and operations when provisioning
virtual machines and cloud instances, and
manage your complete virtual machine lifecycle.
In six parts, this book helps you: Learn the
objects and concepts for developing automation
scripts with CloudForms Automate Customize
the steps and workflows involved in provisioning
virtual machines Create and use service
catalogs, items, dialogs, objects, bundles, and
hierarchies Use CloudForm's updated workflow
to retire and delete virtual machines and
services Orchestrate and coordinate with
external services as part of a workflow Explore
distributed automation processing as well as
argument passing and handling
Preparing for the Certified OpenStack

Administrator Exam - Matt Dorn 2017-08-24
Master the objectives required to pass the Certified OpenStack Administrator exam. About This Book Focuses on providing a clear, concise strategy so you gain the specific skills required to pass the Certified OpenStack Administrator exam Includes exercises and performance-based tasks to ensure all exam objectives can be completed via the Horizon dashboard and command-line interface Includes a free OpenStack Virtual Appliance to practice the objectives covered throughout the book Includes a practice exam to put your OpenStack skills to the test to prove you have what it takes to conquer the live exam Updated for the 2017 exam featuring OpenStack Newton Who This Book Is For This book is for IT professionals, system administrators, DevOps engineers, and software developers with basic Linux command-line and networking knowledge. It's also a great guide for those interested in an entry-level OpenStack position but have limited real-world OpenStack experience. After passing the exam, Certified OpenStack Administrators will prove they have the required skills for the job. What You Will Learn Manage the Keystone identity service by creating and modifying domains, groups, projects, users, roles, services, endpoints, and quotas. Upload Glance images, launch new Nova instances, and create flavors, key pairs, and snapshots. Discover Neutron tenant and provider networks, security groups, routers, and floating IPs. Manage the Cinder block storage service by creating volumes and attaching them to instances. Create Swift containers and set access control lists to allow read/write access to your objects. Explore Heat orchestration templates and create, list, and update stacks. In Detail This book provides you with a specific strategy to pass the OpenStack Foundation's first professional certification: the Certified OpenStack Administrator. In a recent survey, 78% of respondents said the OpenStack skills shortage had deterred them from adopting OpenStack. Consider this an opportunity to increase employer and customer confidence by proving you have the skills required to administrate real-world OpenStack clouds. You will begin your journey by getting well-versed with the OpenStack environment, understanding the benefits of taking the exam, and installing an

included OpenStack all-in-one virtual appliance so you can work through objectives covered throughout the book. After exploring the basics of the individual services, you will be introduced to strategies to accomplish the exam objectives relevant to Keystone, Glance, Nova, Neutron, Cinder, Swift, Heat, and troubleshooting. Finally, you'll benefit from the special tips section and a practice exam to put your knowledge to the test. By the end of the journey, you will be ready to become a Certified OpenStack Administrator! Style and approach Clear, concise, and straightforward with supporting diagrams and lab environment tutorials, this book will help you confidently pass Certified OpenStack Administrator objectives on the Horizon dashboard and command-line interface.

Mastering OpenStack - Second Edition - Chandan Dutta Chowdhury 2016-11-30
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.

Deploying OpenStack - Ken Pepple 2011-08
OpenStack was created with the audacious goal of being the ubiquitous software choice for building public and private cloud infrastructures. In just over a year, it's become the most talked-about project in open source. This concise book introduces OpenStack's general design and primary software components in detail, and shows you how to start using it to build cloud infrastructures. If you're a developer, technologist, or system administrator familiar with cloud offerings such as Rackspace Cloud or Amazon Web Services, *Deploying OpenStack* shows you how to obtain and deploy OpenStack software in a few controlled scenarios. Learn about OpenStack Compute (known as "Nova"), OpenStack Object Store ("Swift"), and OpenStack Image Service

("Glance") Understand common pitfalls in architecting, deploying, and implementing your cloud infrastructure with OpenStack Determine which version of the OpenStack code base best suits your deployment needs Define your deployment scenario and finalize key design choices Install Nova on a single node with either the StackOps distro or an Ubuntu package Be familiar with important configuration options and important administrative commands
[IBM Cloud Private System Administrator's Guide](#) - Ahmed Azraq 2019-05-13

IBM® Cloud Private is an application platform for developing and managing containerized applications across hybrid cloud environments, on-premises and public clouds. It is an integrated environment for managing containers that includes the container orchestrator Kubernetes, a private image registry, a management console, and monitoring frameworks. This IBM Redbooks covers tasks performed by IBM Cloud Private system administrators such as installation for high availability, configuration, backup and restore, using persistent volumes, networking, security, logging and monitoring. Istio integration, troubleshooting and so on. As part of this project we also developed several code examples and you can download those from the IBM Redbooks GitHub location:

<https://github.com/IBMRedbooks>. The authors team has many years of experience in implementing IBM Cloud Private and other cloud solutions in production environments, so throughout this document we took the approach of providing you the recommended practices in those areas. If you are an IBM Cloud Private system administrator, this book is for you. If you are developing applications on IBM Cloud Private, you can see the IBM Redbooks publication *IBM Cloud Private Application Developer's Guide*, SG24-8441.

Ceph: Designing and Implementing Scalable Storage Systems - Michael Hackett 2019-01-31
Get to grips with the unified, highly scalable distributed storage system and learn how to design and implement it. Key Features Explore Ceph's architecture in detail Implement a Ceph cluster successfully and gain deep insights into its best practices Leverage the advanced features of Ceph, including erasure coding,

tiering, and BlueStore Book Description This Learning Path takes you through the basics of Ceph all the way to gaining in-depth understanding of its advanced features. You'll gather skills to plan, deploy, and manage your Ceph cluster. After an introduction to the Ceph architecture and its core projects, you'll be able to set up a Ceph cluster and learn how to monitor its health, improve its performance, and troubleshoot any issues. By following the step-by-step approach of this Learning Path, you'll learn how Ceph integrates with OpenStack, Glance, Manila, Swift, and Cinder. With knowledge of federated architecture and CephFS, you'll use Calamari and VSM to monitor the Ceph environment. In the upcoming chapters, you'll study the key areas of Ceph, including BlueStore, erasure coding, and cache tiering. More specifically, you'll discover what they can do for your storage system. In the concluding chapters, you will develop applications that use Librados and distributed computations with shared object classes, and see how Ceph and its supporting infrastructure can be optimized. By the end of this Learning Path, you'll have the practical knowledge of operating Ceph in a production environment. This Learning Path includes content from the following Packt products: Ceph Cookbook by Michael Hackett, Vikhyat Umrao and Karan Singh Mastering Ceph by Nick Fisk Learning Ceph, Second Edition by Anthony D'Atri, Vaibhav Bhembre and Karan Singh What you will learn Understand the benefits of using Ceph as a storage solution Combine Ceph with OpenStack, Cinder, Glance, and Nova components Set up a test cluster with Ansible and virtual machine with VirtualBox Develop solutions with Librados and shared object classes Configure BlueStore and see its interaction with other configurations Tune, monitor, and recover storage systems effectively Build an erasure-coded pool by selecting intelligent parameters Who this book is for If you are a developer, system administrator, storage professional, or cloud engineer who wants to understand how to deploy a Ceph cluster, this Learning Path is ideal for you. It will help you discover ways in which Ceph features can solve your data storage problems. Basic knowledge of storage systems and GNU/Linux will be

beneficial.

Mastering Ceph - Nick Fisk 2019-03-05

Discover the unified, distributed storage system and improve the performance of applications Key Features Explore the latest features of Ceph's Mimic release Get to grips with advanced disaster and recovery practices for your storage Harness the power of Reliable Autonomic Distributed Object Store (RADOS) to help you optimize storage systems Book Description Ceph is an open source distributed storage system that is scalable to Exabyte deployments. This second edition of Mastering Ceph takes you a step closer to becoming an expert on Ceph. You'll get started by understanding the design goals and planning steps that should be undertaken to ensure successful deployments. In the next sections, you'll be guided through setting up and deploying the Ceph cluster with the help of orchestration tools. This will allow you to witness Ceph's scalability, erasure coding (data protective) mechanism, and automated data backup features on multiple servers. You'll then discover more about the key areas of Ceph including BlueStore, erasure coding and cache tiering with the help of examples. Next, you'll also learn some of the ways to export Ceph into non-native environments and understand some of the pitfalls that you may encounter. The book features a section on tuning that will take you through the process of optimizing both Ceph and its supporting infrastructure. You'll also learn to develop applications, which use Librados and distributed computations with shared object classes. Toward the concluding chapters, you'll learn to troubleshoot issues and handle various scenarios where Ceph is not likely to recover on its own. By the end of this book, you'll be able to master storage management with Ceph and generate solutions for managing your infrastructure. What you will learn Plan, design and deploy a Ceph cluster Get well-versed with different features and storage methods Carry out regular maintenance and daily operations with ease Tune Ceph for improved ROI and performance Recover Ceph from a range of issues Upgrade clusters to BlueStore Who this book is for If you are a storage professional, system administrator, or cloud engineer looking for guidance on building powerful storage solutions for your cloud and on-premise

infrastructure, this book is for you.

[Using the IBM Block Storage CSI Driver in a Red Hat OpenShift Environment](#) - Detlef Helmbrecht
2021-05-10

Red Hat OpenShift container platform is one of the leading enterprise-grade container orchestration platforms. It is designed for rapid deployment of web applications, databases, and microservices. Categorized as a container orchestration Platform as a Service (PaaS), it is based on open industry standards, such as the Container Runtime Interface - Open (CRI-O) and Kubernetes. OpenShift allow developers to focus on the code, while the platform manages the complex IT operations and processes. Although open-source, community-driven container orchestration platforms are available, such as OKD and Kubernetes, this IBM® Redpaper® publication focuses on Red Hat OpenShift. It describes the basic concepts of OpenShift persistent storage architecture and its integration into IBM Cloud® Paks. The deployment of the IBM block storage CSI driver also is discussed. This publication also describes the concepts, technology and current working practices for installing the Container Storage Interface (CSI) plug-in for Kubernetes to use IBM Enterprise Storage platforms for persistent storage coupled with Red Hat OpenShift Container Platform (OCP). This publication also provides an overview of containers, Kubernetes, and OpenShift for context (it is expected that the reader has a working knowledge of these underlying technologies). It also includes architectural examples of the orchestration platform will be given. This paper serves as a guide about how to deploy the CSI driver for block storage by using the DS8000® and Spectrum Virtualize platforms as persistent storage in a Red Hat OpenShift platform. The publication is intended for storage administrators, IT architects, OpenShift technical specialists and anyone who wants to integrate IBM Enterprise storage on OpenShift V4.3/4.4/4.5 on IBM Power, IBM Z®, and x86 systems.

Linux System Programming - Robert Love
2013-05-14

UNIX, UNIX LINUX & UNIX TCL/TK. Write software that makes the most effective use of the Linux system, including the kernel and core

system libraries. The majority of both Unix and Linux code is still written at the system level, and this book helps you focus on everything above the kernel, where applications such as Apache, bash, cp, vim, Emacs, gcc, gdb, glibc, ls, mv, and X exist. Written primarily for engineers looking to program at the low level, this updated edition of Linux System Programming gives you an understanding of core internals that makes for better code, no matter where it appears in the stack. -- Provided by publisher.

OpenShift for Developers - Grant Shipley
2016-08-04

Keen to build web applications for the cloud? Get a quick hands-on introduction to OpenShift, the open source Platform as a Service (PaaS) offering from Red Hat. With this practical guide, you'll learn the steps necessary to build, deploy, and host a complete real-world application on OpenShift without having to slog through long, detailed explanations of the technologies involved. OpenShift enables you to use Docker application containers and the Kubernetes cluster manager to automate the way you create, ship, and run applications. Through the course of the book, you'll learn how to use OpenShift and the Wildfly application server to build and then immediately deploy a Java application online. Learn about OpenShift's core technology, including Docker-based containers and Kubernetes Use a virtual machine with OpenShift installed and configured on your local environment Create and deploy your first application on the OpenShift platform Add language runtime dependencies and connect to a database Trigger an automatic rebuild and redeployment when you push changes to the repository Get a working environment up in minutes with application templates Use commands to check and debug your application Create and build Docker-based images for your application

Red Hat RHCSA 8 Cert Guide - Sander van Vugt
2019-11-04

This is the eBook version of the print title. Learn, prepare, and practice for Red Hat RHCSA 8 (EX200) exam success with this Cert Guide from Pearson IT Certification, a leader in IT Certification learning. Master Red Hat RHCSA 8 EX200 exam topics Assess your knowledge with chapter-ending quizzes Review key concepts

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

Docker Orchestration - Randall Smith
2017-01-24

A concise, fast-paced guide to orchestrating and deploying scalable services with Docker About This Book Explore the new features added to the core Docker Engine to make multi-container orchestration easy Leverage tools such as Docker Machine, Swarm, Compose, and third-party tools such as Kubernetes, Mesosphere, and

CoreOS to orchestrate containers Use Docker Compose with Swarm and apply rolling updates for zero downtime deployments Who This Book Is For This book is aimed at Sysadmins and DevOps engineers who know what Docker does and are now looking to manage multiple containers on multiple hosts using the orchestration feature. What You Will Learn Build scalable, reliable services with Docker See how to manage a service in Docker using Docker Swarm, Kubernetes, and Mesosphere Discover simpler orchestration tools such as CoreOS/Fleet and Rancher Cattle Understand cluster-wide logging, system monitoring, and troubleshooting Build, test, and deploy containers using Continuous Integration Deploy cluster hosts on cloud services and automate your infrastructure In Detail Docker orchestration is what you need when transitioning from deploying containers individually on a single host to deploying complex multi-container apps on many machines. This book covers the new orchestration features of Docker 1.12 and helps you efficiently build, test, and deploy your application using Docker. You will be shown how to build multi-container applications using Docker Compose. You will also be introduced to the building blocks for multi-host Docker clusters such as registry, overlay networks, and shared storage using practical examples. This book gives an overview of core tools such as Docker Machine, Swarm, and Compose which will enhance your orchestration skills. You'll learn how to set up a swarm using the decentralized building block. Next, you'll be shown how to make the most out of the in-built orchestration feature of Docker engine and you'll use third-party tools such as Kubernetes, Mesosphere, and CoreOS to orchestrate your existing process. Finally, you will learn to deploy cluster hosts on cloud services and automate your infrastructure. Style and approach This comprehensive guide will take you through the orchestration feature of Docker. Using practical examples, you will discover various tools that can be used to manage multiple containers with ease.

Kubernetes Operators - Jason Dobies
2020-02-21

Operators are a way of packaging, deploying, and managing Kubernetes applications. A

Kubernetes application doesn't just run on Kubernetes; it's composed and managed in Kubernetes terms. Operators add application-specific operational knowledge to a Kubernetes cluster, making it easier to automate complex, stateful applications and to augment the platform. Operators can coordinate application upgrades seamlessly, react to failures automatically, and streamline repetitive maintenance like backups. Think of Operators as site reliability engineers in software. They work by extending the Kubernetes control plane and API, helping systems integrators, cluster administrators, and application developers reliably deploy and manage key services and components. Using real-world examples, authors Jason Dobies and Joshua Wood demonstrate how to use Operators today and how to create Operators for your applications with the Operator Framework and SDK. Learn how to establish a Kubernetes cluster and deploy an Operator

Examine a range of Operators from usage to implementation

Explore the three pillars of the Operator Framework: the Operator SDK, the Operator Lifecycle Manager, and Operator Metering

Build Operators from the ground up using the Operator SDK

Build, package, and run an Operator in development, testing, and production phases

Learn how to distribute your Operator for installation on Kubernetes clusters

Kubernetes - A Complete DevOps Cookbook

- Murat Karslioglu 2020-03-13

Leverage Kubernetes and container architecture to successfully run production-ready workloads

Key Features

- Implement Kubernetes to orchestrate and scale applications

- Proficiently leverage the latest features of Kubernetes to resolve common as well as complex problems in a cloud-native environment

- Gain hands-on experience in securing, monitoring, and troubleshooting your application

Book Description Kubernetes is a popular open source orchestration platform for managing containers in a cluster environment. With this Kubernetes cookbook, you'll learn how to implement Kubernetes using a recipe-based approach. The book will prepare you to create highly available Kubernetes clusters on multiple clouds such as Amazon Web Services (AWS), Google Cloud Platform (GCP), Azure, Alibaba,

and on-premises data centers. Starting with recipes for installing and configuring Kubernetes instances, you'll discover how to work with Kubernetes clients, services, and key metadata. You'll then learn how to build continuous integration/continuous delivery (CI/CD) pipelines for your applications, and understand various methods to manage containers. As you advance, you'll delve into Kubernetes' integration with Docker and Jenkins, and even perform a batch process and configure data volumes. You'll get to grips with methods for scaling, security, monitoring, logging, and troubleshooting. Additionally, this book will take you through the latest updates in Kubernetes, including volume snapshots, creating high availability clusters with kops, running workload operators, new inclusions around kubectl and more. By the end of this book, you'll have developed the skills required to implement Kubernetes in production and manage containers proficiently. What you will learn

Deploy cloud-native applications on Kubernetes

Automate testing in the DevOps workflow

Discover and troubleshoot common storage issues

Dynamically scale containerized services to manage fluctuating traffic needs

Understand how to monitor your containerized DevOps environment

Build DevSecOps into CI/CD pipelines

Who this book is for

This Kubernetes book is for developers, IT professionals, and DevOps engineers and teams who want to use Kubernetes to manage, scale, and orchestrate applications in their organization. Basic understanding of Kubernetes and containerization is necessary.

How Open Source Ate Software - Gordon Haff 2018-08-21

Learn how free software became open source and how you can sell open source software. This book provides a historical context of how open source has thoroughly transformed how we write software, how we cooperate, how we communicate, how we organize, and, ultimately, how we think about business values. You'll look at project and community examples including Linux, BSD, Apache, and Kubernetes, understand the open source development model, and how open source has influenced approaches more broadly, even proprietary software, such as open betas. You'll also examine the flipside, the "Second Machine Age," and the challenges

of open source-based business models. Today, open source serves as shorthand for much broader trends and behaviors. It's not just about a free (in all senses of the word) alternative to commercial software. It increasingly is the new commercial software. How Open Source Ate Software reveals how open source has much in common, and is often closely allied, with many other trends in business and society. You'll see how it enables projects that go beyond any individual company. That makes open source not just a story about software, but a story about almost everything. What You'll Learn Understand open source opportunities and challenges Sell software if you're giving it away Apply open source principles more broadly to openorg, devops, etc. Review which organizational incentives you can implement Who This Book Is For Anyone who has an interest in what is happening in open source and the open source community, and anyone who is contemplating making a business that involves open source.

Ceph Cookbook - Karan Singh 2016-02-29

Over 100 effective recipes to help you design, implement, and manage the software-defined and massively scalable Ceph storage system About This Book Implement a Ceph cluster successfully and gain deep insights into its best practices Harness the abilities of experienced storage administrators and architects, and run your own software-defined storage system This comprehensive, step-by-step guide will show you how to build and manage Ceph storage in production environment Who This Book Is For This book is aimed at storage and cloud system engineers, system administrators, and technical architects who are interested in building software-defined storage solutions to power their cloud and virtual infrastructure. If you have basic knowledge of GNU/Linux and storage systems, with no experience of software defined storage solutions and Ceph, but eager to learn this book is for you. What You Will Learn Understand, install, configure, and manage the Ceph storage system Get to grips with performance tuning and benchmarking, and gain practical tips to run Ceph in production Integrate Ceph with OpenStack Cinder, Glance, and nova components Deep dive into Ceph object storage, including s3, swift, and keystone

integration Build a Dropbox-like file sync and share service and Ceph federated gateway setup Gain hands-on experience with Calamari and VSM for cluster monitoring Familiarize yourself with Ceph operations such as maintenance, monitoring, and troubleshooting Understand advanced topics including erasure coding, CRUSH map, cache pool, and system maintenance In Detail Ceph is a unified, distributed storage system designed for excellent performance, reliability, and scalability. This cutting-edge technology has been transforming the storage industry, and is evolving rapidly as a leader in software-defined storage space, extending full support to cloud platforms such as Openstack and Cloudstack, including virtualization platforms. It is the most popular storage backend for Openstack, public, and private clouds, so is the first choice for a storage solution. Ceph is backed by RedHat and is developed by a thriving open source community of individual developers as well as several companies across the globe. This book takes you from a basic knowledge of Ceph to an expert understanding of the most advanced features, walking you through building up a production-grade Ceph storage cluster and helping you develop all the skills you need to plan, deploy, and effectively manage your Ceph cluster. Beginning with the basics, you'll create a Ceph cluster, followed by block, object, and file storage provisioning. Next, you'll get a step-by-step tutorial on integrating it with OpenStack and building a Dropbox-like object storage solution. We'll also take a look at federated architecture and CephFS, and you'll dive into Calamari and VSM for monitoring the Ceph environment. You'll develop expert knowledge on troubleshooting and benchmarking your Ceph storage cluster. Finally, you'll get to grips with the best practices to operate Ceph in a production environment. Style and approach This step-by-step guide is filled with practical tutorials, making complex scenarios easy to understand.

OpenStack Object Storage (Swift) Essentials

- Amar Kapadia 2015-05-30

If you are an IT administrator and you want to enter the world of cloud storage using OpenStack Swift, then this book is ideal for you. Basic knowledge of Linux and server technology

is beneficial to get the most out of the book.

OpenShift OKD on IBM LinuxONE, Installation Guide - Narjisse Zaki 2019-10-21

This document provides the step-by-step instructions for installing OpenShift OKD 3.10 on LinuxONE. The intended audience is Systems Architects and Specialists who design, size, and implement solutions on IBM® infrastructures.

Accelerating Modernization with Agile Integration - Adeline SE Chun 2020-07-01

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

Learning Ceph - Second Edition - Anthony D'Atri 2017-10-13

Implement and manage your software-defined, massively scalable storage system
About This Book* Explore Ceph's architecture in order to achieve scalability and high availability* Learn

to utilize Ceph efficiently with the help of practical examples* Successfully implement Ceph clusters to scale-out storage solutions along with outstanding data protection
Who This Book Is ForA basic knowledge of GNU/Linux, and storage systems, and server components is assumed. If you have no experience of software-defined storage solutions and Ceph, but are eager to learn about them, this is the book for you.
What You Will Learn* The limitations of existing systems and why you should use Ceph as a storage solution* Familiarity with Ceph's architecture, components, and services* Instant deployment and testing of Ceph within a Vagrant and VirtualBox environment* Ceph operations including maintenance, monitoring, and troubleshooting* Storage provisioning of Ceph's block, object, and filesystem services* Integrate Ceph with OpenStack* Advanced topics including erasure coding, CRUSH maps, and performance tuning* Best practices for your Ceph clusters
In Detail
Learning Ceph, Second Edition will give you all the skills you need to plan, deploy, and effectively manage your Ceph cluster. You will begin with the first module, where you will be introduced to Ceph use cases, its architecture, and core projects. In the next module, you will learn to set up a test cluster, using Ceph clusters and hardware selection. After you have learned to use Ceph clusters, the next module will teach you how to monitor cluster health, improve performance, and troubleshoot any issues that arise. In the last module, you will learn to integrate Ceph with other tools such as OpenStack, Glance, Manila, Swift, and Cinder. By the end of the book you will have learned to use Ceph effectively for your data storage requirements.
Style and approach
This step-by-step guide, including use cases and examples, not only helps you to easily use Ceph but also demonstrates how you can use it to solve any of your server or drive storage issues.

The Open Organization - Jim Whitehurst 2015

This is a story of reinvention. Jim Whitehurst, celebrated president and CEO of one of the world's most revolutionary software companies, tells first-hand his journey from traditional manager (Delta Air Lines, Boston Consulting Group) and "chief" problem solver to CEO of one of the most open organizational environments

he'd ever encountered. This challenging transition, and what Whitehurst learned in the interim, has paved the way for a new way of managing—one this modern leader sees as the only way companies will successfully function in the future. Whitehurst says beyond embracing the technology that has so far disrupted entire industries, companies must now adapt their management and organizational design to better fit the Information Age. His mantra? “Adapt or die.” Indeed, the successful company Whitehurst leads—the open source giant Red Hat—has become the organizational poster child for how to reboot, redesign, and reinvent an organization for a decentralized, digital age. Based on open source principles of transparency, participation, and collaboration, “open management” challenges conventional business ideas about what companies are, how they run, and how they make money. This book provides the blueprint for putting it into practice in your own firm. He covers challenges that have been missing from the conversation to date, among them: how to scale engagement; how to have healthy debates that net progress; and how to attract and keep the “Social Generation” of workers. Through a mix of vibrant stories, candid lessons, and tested processes, Whitehurst shows how Red Hat has blown the traditional operating model to pieces by emerging out of a pure bottom up culture and learning how to execute it at scale. And he explains what other companies are, and need to be doing to bring this open style into all facets of the organization. By showing how to apply open source methods to everything from structure, management, and strategy to a firm's customer and partner relationships, leaders and teams will now have the tools needed to reach a new level of work. And with that new level of work comes unparalleled success. The Open Organization is your new resource for doing business differently. Get ready to make traditional management thinking obsolete.

Mastering KVM Virtualization - Humble Devassy Chiramal 2016-08-19

Dive in to the cutting edge techniques of Linux KVM virtualization, and build the virtualization solutions your datacentre demands About This Book Become an expert in Linux virtualization Migrate your virtualized datacenter to the cloud Find out how to build a large scale virtualization

solution that will transform your organization Who This Book Is For Linux administrators - if you want to build incredible, yet manageable virtualization solutions with KVM this is the book to get you there. It will help you apply what you already know to some tricky virtualization tasks. What You Will Learn Explore the ecosystem of tools that support Linux virtualization Find out why KVM offers you a smarter way to unlock the potential of virtualization Implement KVM virtualization using oVirt Explore the KVM architecture - so you can manage, scale and optimize it with ease Migrate your virtualized datacenter to the cloud for truly resource-efficient computing Find out how to integrate OpenStack with KVM to take full control of the cloud In Detail A robust datacenter is essential for any organization - but you don't want to waste resources. With KVM you can virtualize your datacenter, transforming a Linux operating system into a powerful hypervisor that allows you to manage multiple OS with minimal fuss. This book doesn't just show you how to virtualize with KVM - it shows you how to do it well. Written to make you an expert on KVM, you'll learn to manage the three essential pillars of scalability, performance and security - as well as some useful integrations with cloud services such as OpenStack. From the fundamentals of setting up a standalone KVM virtualization platform, and the best tools to harness it effectively, including virt-manager, and kimchi-project, everything you do is built around making KVM work for you in the real-world, helping you to interact and customize it as you need it. With further guidance on performance optimization for Microsoft Windows and RHEL virtual machines, as well as proven strategies for backup and disaster recovery, you'll can be confident that your virtualized data center is working for your organization - not hampering it. Finally, the book will empower you to unlock the full potential of cloud through KVM. Migrating your physical machines to the cloud can be challenging, but once you've mastered KVM, it's a little easier. Style and approach Combining advanced insights with practical solutions, Mastering KVM Virtualization is a vital resource for anyone that believes in the power of virtualization to help a business use resources more effectively.

Migrating Linux to Microsoft Azure - Rithin

Skaria 2021-07-28

Discover expert guidance for moving on-premises virtual machines running on Linux servers to Azure by implementing best practices and optimizing costs. Key Features: Work with real-life migrations to understand the dos and don'ts of the process. Deploy a new Linux virtual machine and perform automation and configuration management. Get to grips with debugging your system and collecting error logs with the help of hands-on examples. Book Description: With cloud adoption at the core of digital transformation for organizations, there has been a significant demand for deploying and hosting enterprise business workloads in the cloud. *Migrating Linux to Microsoft Azure* offers a wealth of actionable insights into deploying Linux workload to Azure. You'll begin by learning about the history of IT, operating systems, Unix, Linux, and Windows before moving on to look at the cloud and what things were like before virtualization. This will help anyone new to Linux become familiar with the terms used throughout the book. You'll then explore popular Linux distributions, including RHEL 7, RHEL 8, SLES, Ubuntu Pro, CentOS 7, and more. As you progress, you'll cover the technical details of Linux workloads such as LAMP, Java, and SAP, and understand how to assess your current environment and prepare for your migration to Azure through cloud governance and operations planning. Finally, you'll go through the execution of a real-world migration project and learn how to analyze and debug some common problems that Linux on Azure users may encounter. By the end of this Linux book, you'll be proficient at performing an effective migration of Linux workloads to Azure for your organization. What you will learn: Grasp the terminology and technology of various Linux distributions. Understand the technical support co-operation between Microsoft and commercial Linux vendors. Assess current workloads by using Azure Migrate. Plan cloud governance and operations. Execute a real-world migration project. Manage project, staffing, and customer engagement. Who this book is for: This book is for cloud architects, cloud solution providers, and any stakeholders dealing with migration of Linux workload to Azure. Basic familiarity with

Microsoft Azure would be a plus.

Red Hat and IT Security - Rithik Chatterjee

2020-11-21

Use Red Hat's security tools to establish a set of security strategies that work together to help protect your digital data. You will begin with the basic concepts of IT security and DevOps with topics such as CIA triage, security standards, network and system security controls and configuration, hybrid cloud infrastructure security, and the CI/CD process. Next, you will integrate and automate security into the DevOps cycle, infrastructure, and security as code. You will also learn how to automate with Red Hat Ansible Automation Platform and about hybrid cloud infrastructure. The later chapters will cover hyper-converged infrastructure and its security, Red Hat Smart Management, predictive analytics with Red Hat Insights, and Red Hat security auditing to ensure best security practices. Lastly, you will see the different types of case studies with real-world examples. *Red Hat and IT Security* will help you get a better understanding of IT security concepts from a network and system administration perspective. It will help you to understand how the IT infrastructure landscape can change by implementing specific security best practices and integrating Red Hat products and solutions to counter against modern cybersecurity threats. What You Will Learn ● Understand IT infrastructure security and its best practices ● Implement hybrid cloud infrastructure ● Realign DevOps process into DevSecOps, emphasizing security ● Implement automation in IT infrastructure services using Red Hat Ansible ● Explore Red Hat Smart Management, predictive analytics, and auditing Who This Book Is For: IT professionals handling network/system administration or the IT infrastructure of an organization. DevOps professionals and cybersecurity analysts would find the book useful.

Hyperconverged Infrastructure Data Centers -

Sam Halabi 2019-01-18

Improve Manageability, Flexibility, Scalability, and Control with Hyperconverged Infrastructure. Hyperconverged infrastructure (HCI) combines storage, compute, and networking in one unified system, managed locally or from the cloud. With HCI, you can leverage the cloud's simplicity,

flexibility, and scalability without losing control or compromising your ability to scale. In *Hyperconverged Infrastructure Data Centers*, best-selling author Sam Halabi demystifies HCI technology, outlines its use cases, and compares solutions from a vendor-neutral perspective. He guides you through evaluation, planning, implementation, and management, helping you decide where HCI makes sense, and how to migrate legacy data centers without disrupting production systems. The author brings together all the HCI knowledge technical professionals and IT managers need, whether their background is in storage, compute, virtualization, switching/routing, automation, or public cloud platforms. He explores leading solutions including the Cisco HyperFlex platform, VMware vSAN, Nutanix Enterprise Cloud, Cisco Application-Centric Infrastructure (ACI), VMware's NSX, the open source OpenStack and Open vSwitch (OVS) / Open Virtual Network (OVN), and Cisco CloudCenter for multicloud management. As you explore discussions of automation, policy management, and other key HCI capabilities, you'll discover powerful new opportunities to improve control, security, agility, and performance. Understand and overcome key limits of traditional data center designs Discover improvements made possible by advances in compute, bus interconnect, virtualization, and software-defined storage Simplify rollouts, management, and integration with converged infrastructure (CI) based on the Cisco Unified Computing System (UCS) Explore HCI functionality, advanced capabilities, and benefits Evaluate key HCI applications, including DevOps, virtual desktops, ROBO, edge computing, Tier 1 enterprise applications, backup, and disaster recovery Simplify application deployment and policy setting by implementing a new model for provisioning, deployment, and management Plan, integrate, deploy, provision, manage, and optimize the Cisco HyperFlex hyperconverged infrastructure platform Assess alternatives such as VMware vSAN, Nutanix, open source OpenStack, and OVS/OVN, and compare architectural differences with HyperFlex Compare Cisco ACI (Application-Centric Infrastructure) and VMware NSX approaches to network automation, policies, and security 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.

IBM Spectrum Discover: Metadata Management for Deep Insight of Unstructured Storage - Joseph Dain
2019-10-01

This IBM® Redpaper publication provides a comprehensive overview of the IBM Spectrum® Discover metadata management software platform. We give a detailed explanation of how the product creates, collects, and analyzes metadata. Several in-depth use cases are used that show examples of analytics, governance, and optimization. We also provide step-by-step information to install and set up the IBM Spectrum Discover trial environment. More than 80% of all data that is collected by organizations is not in a standard relational database. Instead, it is trapped in unstructured documents, social media posts, machine logs, and so on. Many organizations face significant challenges to manage this deluge of unstructured data such as: Pinpointing and activating relevant data for large-scale analytics Lacking the fine-grained visibility that is needed to map data to business priorities Removing redundant, obsolete, and trivial (ROT) data Identifying and classifying sensitive data IBM Spectrum Discover is a modern metadata management software that provides data insight for petabyte-scale file and Object Storage, storage on premises, and in the cloud. This software enables organizations to make better business decisions and gain and maintain a competitive advantage. IBM Spectrum Discover provides a rich metadata layer that enables storage administrators, data stewards, and data scientists to efficiently manage, classify, and gain insights from massive amounts of unstructured data. It improves storage economics, helps mitigate risk, and accelerates large-scale analytics to create competitive advantage and speed critical research.

OpenStack Essentials - Dan Radez 2016-08-31
Untangle the complexity of OpenStack clouds through this practical tutorial About This Book
Navigate through the complex jungle of

components in OpenStack using practical instructions This book helps administrators, cloud engineers, and even developers to consolidate and control pools of compute, networking, and storage resources Learn to use the centralized dashboard and administration panel to monitor large-scale deployments Who This Book Is For This book is perfect for administrators, cloud engineers, and operators who want to get started with OpenStack, solve basic problems encountered during deployment, and get up to speed with the latest release of OpenStack. Familiarity with the Linux command line and experience with Linux system administration is expected. What You Will Learn Brush up on the latest release, and how it affects the various components Install OpenStack using the Packstack and RDO Manager installation tool Learn to convert a computer node that supports Docker containers Implement Ceph Block Device images with OpenStack Create and allocate virtual networks, routers and IP addresses to OpenStack Tenants. Configuring and Launching a Docker container. In Detail OpenStack is a widely popular platform for cloud computing. Applications that are built for this platform are resilient to failure and convenient to scale. This book, an update to our extremely popular OpenStack Essentials (published in May 2015) will help you master not only the essential bits, but will also examine the new features of the latest OpenStack release - Mitaka; showcasing how to put them to work straight away. This book begins with the installation and demonstration of the architecture. This book will teach you the core 8 topics of OpenStack. They are Keystone for Identity Management, Glance for Image management, Neutron for network management, Nova for instance management, Cinder for Block storage, Swift for Object storage, Ceilometer for Telemetry and Heat for Orchestration. Further more you will learn about launching and configuring Docker containers and also about scaling them horizontally. You will also learn about monitoring and Troubleshooting OpenStack. Style and approach This book offers step-by-step practical instructions to help you quickly navigate through the complexities of OpenStack

IBM Spectrum Scale CSI Driver for Container Persistent Storage - Abhishek Jain

2020-04-10

IBM® Spectrum Scale is a proven, scalable, high-performance data and file management solution. It provides world-class storage management with extreme scalability, flash accelerated performance, automatic policy-based storage that has tiers of flash through disk to tape. It also provides support for various protocols, such as NFS, SMB, Object, HDFS, and iSCSI. Containers can leverage the performance, information lifecycle management (ILM), scalability, and multisite data management to give the full flexibility on storage as they experience on the runtime. Container adoption is increasing in all industries, and they sprawl across multiple nodes on a cluster. The effective management of containers is necessary because their number will probably reach a far greater number than virtual machines today. Kubernetes is the standard container management platform currently being used. Data management is of ultimate importance, and often is forgotten because the first workloads containerized are ephemeral. For data management, many drivers with different specifications were available. A specification named Container Storage Interface (CSI) was created and is now adopted by all major Container Orchestrator Systems available. Although other container orchestration systems exist, Kubernetes became the standard framework for container management. It is a very flexible open source platform used as the base for most cloud providers and software companies' container orchestration systems. Red Hat OpenShift is one of the most reliable enterprise-grade container orchestration systems based on Kubernetes, designed and optimized to easily deploy web applications and services. OpenShift enables developers to focus on the code, while the platform takes care of all of the complex IT operations and processes. This IBM Redbooks® publication describes how the CSI Driver for IBM file storage enables IBM Spectrum® Scale to be used as persistent storage for stateful applications running in Kubernetes clusters. Through the Container Storage Interface Driver for IBM file storage, Kubernetes persistent volumes (PVs) can be provisioned from IBM Spectrum Scale. Therefore, the containers can be used with stateful microservices, such as database

applications (MongoDB, PostgreSQL, and so on).

Heterogeneity, High Performance Computing, Self-Organization and the Cloud

- Theo Lynn 2018-05-18

This book is open access under a CC BY NC ND license. It addresses the most recent developments in cloud computing such as HPC in the Cloud, heterogeneous cloud, self-organising and self-management, and discusses the business implications of cloud computing adoption. Establishing the need for a new architecture for cloud computing, it discusses a novel cloud management and delivery

architecture based on the principles of self-organisation and self-management. This focus shifts the deployment and optimisation effort from the consumer to the software stack running on the cloud infrastructure. It also outlines validation challenges and introduces a novel generalised extensible simulation framework to illustrate the effectiveness, performance and scalability of self-organising and self-managing delivery models on hyperscale cloud infrastructures. It concludes with a number of potential use cases for self-organising, self-managing clouds and the impact on those businesses.