

Openstack Ceph E Le Nuove Architetture Progetti Cloud

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.

This is not an instructional guide, but a practical, scenario-based book which guides you through everything you need to know in a practical manner by letting you build your own cluster. By the end of the book, you will have a fully functional Proxmox cluster setup at your disposal and have the knowledge to replicate virtualization solutions. If you already know what the word "virtualization" means and you are ready to stand out from the crowd equipped with the unique ability to design and implement a rock-solid virtualized network environment using Proxmox, then you have just picked up the only book you will need. Linux system administration experience together with knowledge of networking and virtualization concepts is assumed. This book is also useful if you are already using Proxmox and simply want to master its advanced features.

This book constitutes the refereed post-conference proceedings of the Third International

Symposium on Benchmarking, Measuring, and Optimization, Bench 2020, held virtually in November 2020. The 12 revised full papers and 1 revised short paper presented were carefully reviewed and selected from 28 submissions. The papers are organized in topical sections named: best paper session; data management and storage; supercomputing; benchmarking on GPU; and application and dataset.

Learn in-demand cloud computing skills from industry experts *Deploying and Managing a Cloud Infrastructure* is an excellent resource for IT professionals seeking to tap into the demand for cloud administrators. This book helps prepare candidates for the CompTIA Cloud+ Certification (CV0-001) cloud computing certification exam. Designed for IT professionals with 2-3 years of networking experience, this certification provides validation of your cloud infrastructure knowledge. With over 30 years of combined experience in cloud computing, the author team provides the latest expert perspectives on enterprise-level mobile computing, and covers the most essential topics for building and maintaining cloud-based systems, including:

- Understanding basic cloud-related computing concepts, terminology, and characteristics
- Identifying cloud delivery solutions and deploying new infrastructure
- Managing cloud technologies, services, and networks
- Monitoring hardware and software performance

Featuring real-world examples and interactive exercises, *Deploying and Managing Cloud Infrastructure* delivers practical knowledge you can apply immediately. And, in addition, you also get access to a full set of electronic study tools including:

- Interactive Test Environment
- Electronic Flashcards
- Glossary of Key Terms

Now is the time to learn the cloud computing skills you need to take that next step in your IT career.

OpenStack(r) is the leading open source IaaS platform, powering many of the world's most

notable science and research organisations. Surprisingly, research and science disciplines comprise some of the most prevalent use cases for OpenStack clouds, and OpenStack has provided compelling solutions for many of the challenges of delivering exible infrastructure for research computing. High-performance computing (HPC) and high-throughput computing (HTC) workloads require massive scaling and cluster networking; storage, compute and networking access to large volumes of data; and workload and infrastructure manageability. OpenStack software supports these needs today and the development community is rapidly expanding services to fill gaps. By managing resources as an OpenStack private cloud, researchers are able to work in environments tailored to their requirements. The dynamic, automated nature of software-defined infrastructure cuts away time wasted on the distractions of setup, and enables researchers to maximise the time they spend on research itself. This paper is intended for HPC system architects and research computing managers that are exploring the benefits of cloud and how to bring those benefits to HPC workloads. It can also be used by current OpenStack users to delve into additional capabilities. This is a deep dive into the important functions, considerations, and further reading. Each section includes user examples that describe real-world architecture and operations.

OpenStack is today's leading technology for building and integrating public and private clouds. *Common OpenStack Deployments* is a complete, practical guide to deploying OpenStack and understanding its internals. Key project contributor Elizabeth Joseph, with expert implementer Matt Fischer, shares up-to-date recipes for deploying OpenStack on both virtual and physical servers, and for using OpenStack to address any real-world challenge. First, Joseph and Fischer help you master OpenStack concepts and components by guiding you through small-

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

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 infrastructureWho This Book Is ForThis 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 improvementIn DetailIn 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.

Develop and manage applications on the AWS and OpenStack platforms with this comprehensive learning guide. Key Features A step-by-step guide to help you develop applications on the hybrid cloud platform. Acquire an in-depth understanding of the OpenStack and AWS cloud platforms. Extensive source code examples for OpenStack and AWS applications. Easily troubleshoot OpenStack and AWS issues. Understand the best practices and security measures for the hybrid cloud platform. Book Description This book introduces you to the hybrid cloud platform, and focuses on the AWS public cloud and OpenStack private cloud platforms. It provides a deep dive into the AWS and OpenStack cloud platform services that are essential for developing hybrid cloud applications. You will learn to develop applications on AWS and OpenStack platforms with ease by leveraging various cloud services and taking advantage of PaaS. The book provides you with the ability to leverage the flexibility of choosing a cloud platform for migrating your existing resources to the cloud, as well as developing hybrid cloud applications that can migrate virtual machine instances from AWS to OpenStack and vice versa. You will also be able to build and test cloud applications without worrying about the system that your development environment supports. The book also

provides an in-depth understanding of the best practices that are followed across the industry for developing cloud applications, as well as for adapting the hybrid cloud platform. Lastly, it also sheds light on various troubleshooting techniques for OpenStack and AWS cloud platform services that are consumed by hybrid cloud applications. By the end of this book, you will have a deep understanding of the hybrid cloud platform and will be able to develop robust, efficient, modular, scalable, and flexible cloud applications. What you will learn

- Understand the hybrid cloud platform
- Explore the AWS and OpenStack cloud platforms in depth
- Develop AWS applications with source code examples
- Develop OpenStack applications with source code examples
- Troubleshoot OpenStack and AWS
- Learn hybrid cloud best practices
- Understand security measures on the hybrid cloud

Who this book is for If you are an IT professional, developer, or a DevOps engineer looking to develop and manage your applications on the hybrid cloud platform, then this book is for you. Some prior knowledge of the public and private cloud will enhance your skills. Developers looking to build applications using AWS or OpenStack services will also benefit from this book.

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.

The superabundance of data that is created by today's businesses is making storage a strategic investment priority for companies of all sizes. As storage takes precedence, the following major initiatives emerge: Flatten and converge your network: IBM® takes an open, standards-based approach to implement the latest advances in the flat, converged data center network designs of today. IBM Storage solutions enable clients to deploy a high-speed, low-latency Unified Fabric Architecture. Optimize and automate virtualization: Advanced virtualization awareness reduces the cost and complexity of deploying physical and virtual data center infrastructure. Simplify management: IBM data center networks are easy to deploy, maintain, scale, and virtualize, delivering the foundation of consolidated operations for dynamic infrastructure management. Storage is no longer an afterthought. Too much is at stake. Companies are searching for more ways to efficiently manage expanding volumes of data, and to make that data accessible throughout the enterprise. This demand is propelling the move of storage into the network. Also, the increasing complexity of managing large numbers of storage devices and vast amounts of data is driving greater business value into software and services. With current estimates of the amount of data to be managed and made available increasing at 60% each year, this outlook is where a storage area network (SAN) enters the arena. SANs are the leading storage infrastructure for the global economy of today. SANs offer simplified storage management, scalability, flexibility, and availability; and improved data access, movement, and backup. Welcome to the

cognitive era. The smarter data center with the improved economics of IT can be achieved by connecting servers and storage with a high-speed and intelligent network fabric. A smarter data center that hosts IBM Storage solutions can provide an environment that is smarter, faster, greener, open, and easy to manage. This IBM® Redbooks® publication provides an introduction to SAN and Ethernet networking, and how these networks help to achieve a smarter data center. This book is intended for people who are not very familiar with IT, or who are just starting out in the IT world. This book constitutes the reviewed proceedings of the first Conference on Performance Evaluation and Benchmarking, TPCTC 2009, held in Lyon, France, August 24-28,2009. The 16 full papers and two keynote papers were carefully selected from 34 submissions. This book considers issues such as appliance, business intelligence, cloud computing, complex event processing, database performance optimizations, green computing, data compression, disaster tolerance and recovery, energy and space efficiency, hardware innovations, high speed data generation, hybrid workloads or operational data warehousing, unstructured data management, software management and maintenance, virtualization and very large memory systems

Leverage the best SDN technologies for your OpenStack-based cloud infrastructure
About This Book Learn how to leverage critical SDN technologies for OpenStack
Networking APIs via plugins and drivers Champion the skills of achieving complete
SDN with OpenStack with specific use cases and capabilities only covered in this title

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

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

As workloads are being offloaded to IBM® LinuxONE based cloud environments, it is important to ensure that these workloads and environments are secure. This IBM Redbooks® publication describes the necessary steps to secure your environment from the hardware level through all of the components that are involved in a LinuxONE cloud infrastructure that use Linux and IBM z/VM®. The audience for this book is IT architects, IT Specialists, and those users who plan to use LinuxONE for their cloud environments.

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.

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.

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

Understand the fundamental factors of data storage system performance and master an essential analytical skill using block trace via applications such as MATLAB and Python tools. You will increase your productivity and learn the best techniques for doing specific tasks (such as analyzing the IO pattern in a quantitative way, identifying the storage system bottleneck, and designing the cache policy). In the new era of IoT, big data, and cloud systems, better performance and higher density of storage systems has become crucial. To increase data storage density, new techniques have evolved and hybrid and parallel access techniques—together with specially designed IO scheduling and data migration algorithms—are being deployed to develop high-performance data storage solutions. Among the various storage system performance analysis techniques, IO event trace analysis (block-level trace analysis particularly) is one of the most common approaches for system optimization and design. However, the task of completing a systematic survey is challenging and very few works on this topic exist. Block Trace Analysis and Storage System Optimization brings together theoretical

analysis (such as IO qualitative properties and quantitative metrics) and practical tools (such as trace parsing, analysis, and results reporting perspectives). The book provides content on block-level trace analysis techniques, and includes case studies to illustrate how these techniques and tools can be applied in real applications (such as SSHD, RAID, Hadoop, and Ceph systems). What You'll Learn Understand the fundamental factors of data storage system performance Master an essential analytical skill using block trace via various applications Distinguish how the IO pattern differs in the block level from the file level Know how the sequential HDFS request becomes "fragmented" in final storage devices Perform trace analysis tasks with a tool based on the MATLAB and Python platforms Who This Book Is For IT professionals interested in storage system performance optimization: network administrators, data storage managers, data storage engineers, storage network engineers, systems engineers

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).

This book contains selected papers from the 9th International Conference on Information Science and Applications (ICISA 2018) and provides a snapshot of the latest issues encountered in technical convergence and convergences of security technology. It explores how information science is core to most current research, industrial and commercial activities and consists of contributions covering topics including Ubiquitous Computing, Networks and Information Systems, Multimedia and Visualization, Middleware and Operating Systems, Security and Privacy, Data Mining and Artificial Intelligence, Software Engineering, and Web Technology. The proceedings introduce the most recent information technology and ideas, applications and problems related to technology convergence, illustrated through case studies, and reviews converging existing security techniques. Through this volume, readers will gain an understanding of the current state-of-the-art information strategies and technologies of convergence security. The intended readership includes researchers in academia, industry and other research institutes focusing on information science and technology.

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.

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

This IBM® Redpaper™ publication takes you on a journey that surveys cloud computing to answer several fundamental questions about storage cloud technology. What are storage clouds? How can a storage cloud help solve your current and future data storage business requirements? What can IBM do to help you implement a storage cloud solution that addresses these needs? This paper shows how IBM storage clouds use the extensive cloud computing experience, services, proven technologies, and products of IBM to support a smart storage cloud solution designed for your storage optimization efforts. Clients face many common storage challenges and some have variations that make them unique. It describes various successful client storage cloud implementations and the options that are available to meet your current needs and position you to avoid storage issues in the future. IBM Cloud™ Services (IBM Cloud Managed Services® and IBM SoftLayer®) are highlighted as well as the contributions of IBM to OpenStack cloud storage. This paper is intended for anyone who wants to learn about storage clouds and how IBM addresses data storage challenges with smart storage cloud solutions. It is suitable for IBM clients, storage solution integrators, and IBM specialist sales representatives.

Today, new business models in the marketplace coexist with traditional ones and their well-established IT architectures. They generate new business needs and

new IT requirements that can only be satisfied by new service models and new technological approaches. These changes are reshaping traditional IT concepts. Cloud in its three main variants (Public, Hybrid, and Private) represents the major and most viable answer to those IT requirements, and software-defined infrastructure (SDI) is its major technological enabler. IBM® technology, with its rich and complete set of storage hardware and software products, supports SDI both in an open standard framework and in other vendors' environments. IBM services are able to deliver solutions to the customers with their extensive knowledge of the topic and the experiences gained in partnership with clients. This IBM Redpaper™ publication focuses on software-defined storage (SDS) and IBM Storage Systems product offerings for software-defined environments (SDEs). It also provides use case examples across various industries that cover different client needs, proposed solutions, and results. This paper can help you to understand current organizational capabilities and challenges, and to identify specific business objectives to be achieved by implementing an SDS solution in your enterprise.

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.

This IBM® Redbooks® publication provides advice and technical information about optimizing and tuning application code to run on systems that are based on the IBM POWER7® and POWER7+™ processors. This advice is drawn from application optimization efforts across many different types of code that runs under the IBM AIX® and Linux operating systems, focusing on the more

pervasive performance opportunities that are identified, and how to capitalize on them. The technical information was developed by a set of domain experts at IBM. The focus of this book is to gather the right technical information, and lay out simple guidance for optimizing code performance on the IBM POWER7 and POWER7+ systems that run the AIX or Linux operating systems. This book contains a large amount of straightforward performance optimization that can be performed with minimal effort and without previous experience or in-depth knowledge. This optimization work can:

- Improve the performance of the application that is being optimized for the POWER7 system
- Carry over improvements to systems that are based on related processor chips
- Improve performance on other platforms

The audience of this book is those personnel who are responsible for performing migration and implementation activities on IBM POWER7-based servers, which includes system administrators, system architects, network administrators, information architects, and database administrators (DBAs).

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 the following

examples: 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. This IBM Redbooks® publication presents several use cases that are focused on artificial intelligence (AI) solutions with IBM Spectrum Discover. This book helps storage administrators and technical specialists plan and implement AI solutions by using IBM Spectrum Discover and several other IBM Storage products.

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

The book at hand gives an overview of the state of the art research in Computational Sustainability as well as case studies of different application scenarios. This covers topics such as renewable energy supply, energy storage and e-mobility, efficiency in data centers and networks, sustainable food and water supply, sustainable health, industrial production and quality, etc. The book describes computational methods and possible application scenarios.

Learning CephPackt Publishing Ltd

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.

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.

There's a lot of information about big data technologies, but splicing these technologies into an end-to-end enterprise data platform is a daunting task not widely covered. With this practical book, you'll learn how to build big data infrastructure both on-premises and in the cloud and successfully architect a modern data platform. Ideal for enterprise architects, IT managers, application architects, and data engineers, this book shows you how to overcome the many challenges that emerge during Hadoop projects. You'll explore the vast landscape of tools available in the Hadoop and big data realm in a thorough technical primer before diving into:

- Infrastructure: Look at all component layers in a modern data platform, from the server to the data center, to establish a solid foundation for data in your enterprise
- Platform: Understand aspects of deployment, operation, security, high availability, and disaster recovery, along with everything you need to know to integrate your platform with the rest of your enterprise IT
- Taking Hadoop to the cloud: Learn the important architectural aspects of running a big data platform in the cloud while maintaining enterprise security and high availability

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.

Presents information on getting the most out of a PC's hardware and software, covering such topics as upgrading the BIOS, configuring the hard drive, installing more RAM, improving CPU performance, and adding COM ports.

Software-Defined Data Infrastructures Essentials provides fundamental coverage of physical, cloud, converged, and virtual server storage I/O networking technologies, trends, tools, techniques, and tradecraft skills. From webscale, software-defined, containers, database, key-value store, cloud, and enterprise to small or medium-size business, the book is filled with techniques, and tips to help develop or refine your server storage I/O hardware, software, and services skills. Whether you are new to data

infrastructures or a seasoned pro, you will find this comprehensive reference indispensable for gaining as well as expanding experience with technologies, tools, techniques, and trends. We had a front row seat watching Greg present live in our education workshop seminar sessions for ITC professionals in the Netherlands material that is in this book. We recommend this amazing book to expand your converged and data infrastructure knowledge from beginners to industry veterans. —Gert and Frank Brouwer, Brouwer Storage Consultancy Software-Defined Data Infrastructures Essentials provides the foundational building blocks to improve your craft in several areas including applications, clouds, legacy, and more. IT professionals, as well as sales professionals and support personnel, stand to gain a great deal by reading this book.—Mark McSherry, Oracle Regional Sales Manager Looking to expand your data infrastructure IQ? From CIOS to operations, sales to engineering, this book is a comprehensive reference, a must read for IT infrastructure professionals, beginners to seasoned experts.—Tom Becchetti, Advisory Systems Engineer Greg Schulz has provided a complete ‘toolkit’ for storage management along with the background and framework for the storage or data infrastructure professional or those aspiring to become one.—Greg Brunton, Experienced Storage and Data Management Professional 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.

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

environments is expected.

This book aims to provide a deep look into Italian actions taken in some fields of science and high performance computing (HPC), and the Italian effort to bridge the HPC gap with respect to Europe. The Italian PON ReCaS Project is written for graduate readers and professionals in the field of high performance computing. It presents and discusses innovative and important technological solutions, and describes interesting results in various fields of application.

ReCaS stands for "Rete di Calcolo per SuperB e altre applicazioni" and is a computing network infrastructure in Southern Italy devoted to scientific and non-scientific applications within the vision of a common European infrastructure for computing, storage and network. The ReCaS project is part of the 2007–2013 European Union strategy, and was funded by the Italian Ministry of Research and Education (MIUR) for the development and enhancement of a distributed computing infrastructure of the Grid/Cloud type over the four EU 'Convergence' regions in Southern Italy: Campania, Puglia and Sicily and Calabria. The network will be open and accessible to all researchers, public and private, and will be characterized by unprecedented computing power and storage capacity. Posted in the European Grid Infrastructure EGI, ReCaS is also an opportunity to the countries of the Mediterranean area and extends the potential of the current network.

Get an in-depth tour of OpenShift, the container-based software deployment and management platform from Red Hat that provides a secure multi-tenant environment for the enterprise. This practical guide describes in detail how OpenShift, building on Kubernetes, enables you to automate the way you create, ship, and run applications in a containerized environment.

Author Graham Dumbleton provides the knowledge you need to make the best use of the

OpenShift container platform to deploy not only your cloud-native applications, but also more traditional stateful applications. Developers and administrators will learn how to run, access, and manage containers in OpenShift, including how to orchestrate them at scale. Build application container images from source and deploy them Implement and extend application image builders Use incremental and chained builds to accelerate build times Automate builds by using a webhook to link OpenShift to a Git repository Add configuration and secrets to the container as project resources Make an application visible outside the OpenShift cluster Manage persistent storage inside an OpenShift container Monitor application health and manage the application lifecycle This book is a perfect follow-up to OpenShift for Developers: A Guide for Impatient Beginners (O'Reilly).

Master the art of container management utilizing the power of Kubernetes. About This Book This practical guide demystifies Kubernetes and ensures that your clusters are always available, scalable, and up to date Discover new features such as autoscaling, rolling updates, resource quotas, and cluster size Master the skills of designing and deploying large clusters on various cloud platforms Who This Book Is For The book is for system administrators and developers who have intermediate level of knowledge with Kubernetes and are now waiting to master its advanced features. You should also have basic networking knowledge. This advanced-level book provides a pathway to master Kubernetes. What You Will Learn Architect a robust Kubernetes cluster for long-time operation Discover the advantages of running Kubernetes on GCE, AWS, Azure, and bare metal See the identity model of Kubernetes and options for cluster federation Monitor and troubleshoot Kubernetes clusters and run a highly available Kubernetes Create and configure custom Kubernetes resources and use third-party

resources in your automation workflows Discover the art of running complex stateful applications in your container environment Deliver applications as standard packages In Detail Kubernetes is an open source system to automate the deployment, scaling, and management of containerized applications. If you are running more than just a few containers or want automated management of your containers, you need Kubernetes. This book mainly focuses on the advanced management of Kubernetes clusters. It covers problems that arise when you start using container orchestration in production. We start by giving you an overview of the guiding principles in Kubernetes design and show you the best practises in the fields of security, high availability, and cluster federation. You will discover how to run complex stateful microservices on Kubernetes including advanced features as horizontal pod autoscaling, rolling updates, resource quotas, and persistent storage back ends. Using real-world use cases, we explain the options for network configuration and provides guidelines on how to set up, operate, and troubleshoot various Kubernetes networking plugins. Finally, we cover custom resource development and utilization in automation and maintenance workflows. By the end of this book, you'll know everything you need to know to go from intermediate to advanced level. Style and approach Delving into the design of the Kubernetes platform, the reader will be exposed to the advanced features and best practices of Kubernetes. This book will be an advanced level book which will provide a pathway to master Kubernetes Teaches you how and what to study in order to be best prepared for the Certified OpenStack Administrator exam. This fast-growing technology is creating a market that needs more qualified IT specialists with proven skills. This book covers 100% of the exam requirements for both The OpenStack Foundation and the Mirantis OpenStack Certification Exam. Each theme

Read PDF Openstack Ceph E Le Nuove Architetture Progetti Cloud

is taught using practical exercises and instructions for the command line and for the graphical client (Horizon). Each chapter is followed by review questions, complete with answers. Even after you have taken and passed your OpenStack exam, this book will remain a useful reference. What You Will Learn Understand the components that make up the cloud. Install and make an OpenStack distribution from Mirantis, Red Hat or another community version. Work with OpenStack Identity Management, Dashboard, CLI, Object Storage, Block Storage, Networking, Telemetry, Orchestration, and Image Services. Learn how to troubleshoot all the main OpenStack services. Understand where to find information for future work with OpenStack. Who This Book Is For Certified OpenStack Administrator Study Guide is for Cloud and Linux engineers looking for a better understanding of how to work with the modern OpenStack IaaS Cloud, and wants to prove their knowledge by passing a Certified OpenStack Administrator Exam.

[Copyright: df5a0f7cb93fbf0f19c75512b8a28db0](https://www.pdfdrive.com/openstack-ceph-e-le-nuove-architetture-progetti-cloud-p123456789.html)