

Compliments of  redhat.

Red Hat Special Edition

OpenStack[®] Storage

FOR
DUMMIES[®]
A Wiley Brand

Learn:

- The basics of OpenStack and Ceph[®] storage
- Configuration best practices for OpenStack and Ceph together
- Why Red Hat Ceph Storage is great for your enterprise

Sébastien Han
Sean Cohen
Daniel Gilfix



ABOUT RED HAT

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to reliable and high-performing cloud, Linux®, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT. Learn more at <http://www.redhat.com>.

OpenStack Storage

FOR
DUMMIES[®]
A Wiley Brand

Red Hat[®] Special Edition

**by Sébastien Han, Sean Cohen,
and Daniel Gilfix**

FOR
DUMMIES[®]
A Wiley Brand

OpenStack Storage For Dummies®, Red Hat® Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2016 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and JBoss are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. The OpenStack Word Mark and OpenStack Logo are either registered trademarks / service marks or trademarks / service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Java is the registered trademark of Oracle America, Inc. in the United States and other countries. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-29270-8 (pbk); ISBN 978-1-119-29253-1 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Project Editor: Carrie Burchfield
Acquisitions Editor: Amy Fandrei
Editorial Manager: Rev Mengle

Business Development Representative:
Susan Blessing
Production Editor: Siddique Shaik

Table of Contents

.....

Introduction	1
About This Book	2
Foolish Assumptions	2
Icons Used in This Book.....	3
Beyond the Book.....	4
Where to Go from Here	4
Chapter 1: Understanding the Basics of OpenStack and Ceph Storage	5
OpenStack Storage.....	5
Red Hat Ceph Storage.....	6
Chapter 2: Exploring Use Cases and Configuration Best Practices	11
Common Use Case Examples	11
On-premises private cloud	12
Public cloud.....	12
NFV infrastructure.....	12
Application development	13
Configuration Best Practices	13
Hyperconverged infrastructure.....	14
Multi-region deployments.....	16
Chapter 3: Ten Reasons to Choose Red Hat Ceph Storage	19

Introduction

The rise of cloud computing continues as more and more companies invest in their enterprise IT capabilities to enable innovation at scale. But just in case you believe that the cloud — like the Internet — is just a passing fad, here's some data to convince you otherwise:

- ✓ According to 451 Research, spending on cloud computing infrastructure is projected to grow at a 30 percent compound annual growth rate (CAGR) through 2018, compared to only 5 percent growth for overall enterprise IT.
- ✓ According to Goldman Sachs, cloud models like Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) will double in corporate IT spend percentage by 2018.
- ✓ According to Forbes, 93 percent of organizations surveyed are *already* running applications or experimenting with IaaS, thereby spurring investments in private cloud computing.

As more applications and data are distributed across virtual machines in the cloud, a more agile and scalable solution is needed to enable ever expanding cloud services and, ultimately, the software-defined data center supporting them. OpenStack is the largest and fastest growing open source cloud infrastructure project today — and the preferred cloud and IaaS technology spanning private and hybrid cloud environments.

OpenStack is scale-out technology that needs scale-out storage to succeed. *Ceph* is a software-defined storage technology solution that provides this scale-out capability with an extensible architecture that integrates more tightly with OpenStack than traditional storage solutions. The OpenStack User Survey by the OpenStack Foundation consistently reveals that OpenStack users overwhelmingly prefer Ceph over other

storage alternatives because of its robust features, tight integration, and original design.

The April 2016 OpenStack User Survey also showcases the fact that Ceph block storage dominates OpenStack production deployments.

Red Hat Ceph Storage is an enterprise-ready implementation of Ceph that provides a single platform solution for OpenStack that's open, adaptable, massively scalable, technologically advanced, and supported worldwide. Red Hat Ceph Storage combines innovation from the open source community with the backing of Red Hat engineering, consulting, and support. Red Hat Ceph Storage provides tight integration with OpenStack services and enables support for block, object, and file storage on standard servers and disks — at a lower cost than traditional storage. Enterprises benefit from the efficiency of self-manageability, dynamic data distribution, and a self-healing architecture, as well as the flexibility of user-driven storage lifecycle management with complete API (application programming interface) coverage.

About This Book

This book consists of three short chapters. Chapter 1 covers the basics of OpenStack and Ceph storage concepts and architectures. Chapter 2 describes OpenStack/Ceph real-world use cases and configuration best practices, and Chapter 3 explains why Red Hat Ceph Storage is the best solution to deploy Ceph in your enterprise.

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I assume a few things about you, the reader, nonetheless:

- ✓ I assume you're a cloud administrator, cloud operator, infrastructure architect, IT manager, or storage administrator in need of scalable, unified storage for virtual machines, images, volumes, objects, shared file systems, and/or databases with OpenStack.

- ✔ I also assume that you work in an open source environment and are somewhat familiar with open source technologies, such as OpenStack and Ceph. Thus, I won't waste any valuable real estate in this short book — or any of your valuable time — espousing the advantages of open source solutions or convincing you of the innate superiority of open source over closed source solutions.
- ✔ Lastly, I assume that you are somewhat technical. In fact, “non-technical Linux user” is probably a bit of an oxymoron. As such, this book is written at a somewhat technical level — but I promise it's not as deep or dry as a HOWTO doc.

If these assumptions describe you, then this book is for you. If none of these assumptions describe you, keep reading anyway. It's a great book and when you finish reading it, you'll know a few things about OpenStack and Ceph storage.

Icons Used in This Book

Throughout this book, I occasionally use special icons to call attention to important information. Here's what to expect:



This icon points out information that you should commit to your non-volatile memory, your gray matter, or your noggin' — along with anniversaries and birthdays!



You won't find a map of the human genome here, but if you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon!



Thank you for reading, hope you enjoy the book, please take care of your writers! Seriously, this icon points out helpful suggestions and useful nuggets of information.



This icon points out the stuff your mother warned you about. Okay, probably not. But you should take heed nonetheless — you might just save yourself some time and frustration!

Beyond the Book

There's only so much I can cover in 24 short pages, so if you find yourself at the end of this book, thinking "gosh, this was an amazing book, where can I learn more?" just go to www.redhat.com/storage.

Where to Go from Here

With my apologies to Lewis Carroll, Alice, and the Cheshire cat: "Would you tell me, please, which way I ought to go from here?"

"That depends a good deal on where you want to get to," said the Cat — err, the Dummies Man.

"I don't much care where . . .," said Alice.

"Then it doesn't matter which way you go!"

That's certainly true of *OpenStack Storage For Dummies*, Red Hat Special Edition, which, like *Alice in Wonderland*, is also destined to become a timeless classic!

If you don't know where you're going, any chapter will get you there — but Chapter 1 might be a good place to start! But if you see a particular topic that piques your interest, feel free to jump ahead to that chapter. Each chapter is written to stand on its own, so you can start reading anywhere and skip around to your heart's content! Read this book in any order that suits you (though I don't recommend upside down or backwards).

I promise you won't get lost falling down the rabbit hole!

Chapter 1

Understanding the Basics of OpenStack and Ceph Storage

In This Chapter

- ▶ Getting to know storage options in OpenStack
 - ▶ Learning about the advantages of Ceph
-

In this chapter, you discover the different storage types supported by OpenStack and how Ceph integrates with OpenStack to provide maximum storage performance and flexibility.

OpenStack Storage

Storage is an important component in the OpenStack cloud architecture. OpenStack supports three types of persistent storage:

- ✓ **Object storage** (implemented in OpenStack as *Swift*) uses variable-sized data containers, which are organized into a flat address space and accessed through a REST API (REpresentational State Transfer application programming interface). Object storage can be primarily used to store applications data, as well as the image catalog of your cloud and your virtual machine snapshots. Amazon S3 is one example of an object storage system.

- ✓ **Block (or volume) storage** (implemented in OpenStack as *Cinder*) uses fixed-length data blocks that are attached to running virtual machines. Block storage is used to add persistent storage to a virtual machine.
- ✓ **Shared file system storage** (implemented in OpenStack as *Manila*) implements a distributed file system solution that allows you to expose and consume your data through a file interface. This process allows multiple instances to access the same shared service. It supports multiple back ends in the form of drivers. Share servers are, mostly, virtual machines that export file shares via different protocols such as NFS, CIFS, GlusterFS, or HDFS. Like the *cinder* block storage, the file system storage is also persistent and used to add persistent storage to a virtual machine and detach storage from one instance to another without data loss.



OpenStack also supports a fourth type of storage — ephemeral storage, which is non-persistent storage, typically used by the OpenStack compute service (Nova) to run an operating system and as scratch space. With ephemeral storage, when a user's session ends, the virtual machine and any of its associated data goes poof — you could say it dissipates in the cloud.

Red Hat Ceph Storage

Red Hat Ceph Storage is a massively scalable, open, software-defined storage system that runs on standard hardware and was built from the ground up to deliver next-generation storage for cloud and emerging workloads.

Red Hat Ceph Storage combines Ceph v10.2 (code name: Jewel) for object, block, and file storage with an integrated on-premises management console called Red Hat Storage Console 2, featuring Ansible-based deployment tools, a graphical user interface with cluster visualization, advanced Ceph monitoring and diagnostic information, as well as cluster and per node usage and performance statistics. Along with support services and seamless integration with Red Hat OpenStack Platform, it provides a crucial storage service for virtual machines, volumes, and images, and a fully supported cloud platform.



Red Hat Ceph Storage also integrates with Ubuntu OpenStack.

Large organizations that use OpenStack can serve thousands of clients or more. Each OpenStack client with its own particular needs consumes block storage resources. Deploying Glance (image service catalog), Cinder (block storage service), and/or Nova (compute service) on a single node becomes impossible to manage in large deployments with thousands of clients or more. Scaling OpenStack externally resolves this challenge.

However, it also makes it a practical requirement to virtualize the storage layer with a solution like Red Hat Ceph Storage so you can scale the Red Hat OpenStack Platform storage layer from tens of terabytes to petabytes, or even exabytes of storage. Red Hat Ceph Storage provides this storage virtualization layer with high availability and high performance while running on commodity hardware. While virtualization might seem like it comes with a performance penalty, Ceph stripes block device images as objects across the cluster. This means large Ceph block device images have better performance than a standalone disk. Ceph block devices also support snapshots, caching, copy-on-write cloning, and copy-on-read cloning for enhanced performance.

The Ceph block device (RBD) integrates tightly with OpenStack by providing an excellent single backend for Glance, Cinder, and Nova that lets users efficiently store images and snapshots, volumes and virtual machine root disks (see Figure 1-1). In addition, the Ceph object gateway (RGW) provides a RESTful interface that's compatible with applications written for OpenStack Swift. Both of these data services are based on the underlying RADOS object store, allowing users to deploy and manage them both within a single cluster.

Ceph is particularly well-suited as a storage backend for Red Hat OpenStack Platform, which provides an integrated and optimized foundation for building OpenStack clouds.

Some additional Red Hat Ceph Storage features include

- ✓ **Horizontal scalability:** Scale-out architecture that grows a cluster from one node to thousands, automatic rebalancing using a peer-to-peer architecture to add capacity at any time with minimal operational effort (say goodbye to forklift upgrades and data migration projects), and hot or phased software upgrades that enable cluster upgrades in phases with minimal or no downtime.

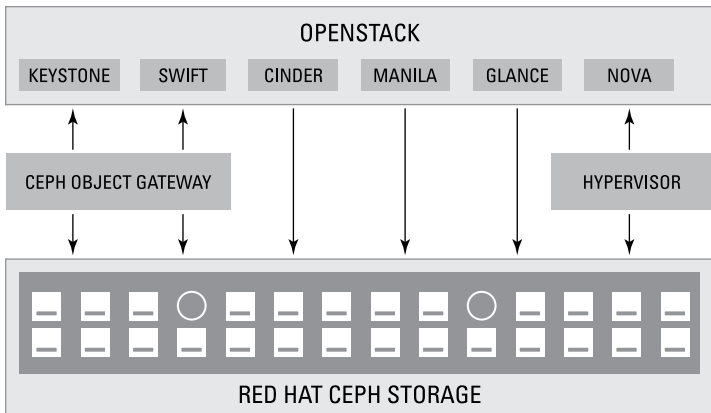


Figure 1-1: Red Hat Ceph Storage architecture.

- ✓ **API support:** Cloud integration with S3 and Swift protocols used by Amazon Web Services and the OpenStack Object Storage project, along with the ability to manage all cluster and object storage functions programmatically.
- ✓ **Security:** Authentication and authorization through Active Directory, LDAP, and Keystone v3, policies for granular control over access at the pool, user, bucket, or data level, and cluster-level at-rest encryption.
- ✓ **Reliability and availability:** Dynamic block resizing expands or shrinks Ceph block devices with zero downtime. Striping, erasure coding, and replication across nodes provides data durability, high availability, and high performance. Storage policies enable users to configure placement of data to reflect service-level agreements (SLAs), performance requirements, and failure domains using Ceph's unique CRUSH algorithm. Snapshots of entire pool or individual block devices are available, and automatic failover prevents server or disk failures from impacting data integrity, availability, or performance.
- ✓ **Performance:** Client-cluster data path for clients to share their I/O model across the entire cluster; copy-on-write cloning to instantly provision tens or hundreds of virtual machine images; in-memory client-side caching to enhance client I/O using a hypervisor cache; and server-side journaling to accelerate the write performance of data by serializing writes.

- ✔ **Multi-data center support and disaster recovery:** Zones and region support enable you to deploy the object storage topologies of Amazon Web Services S3; global clusters enables creating a global namespace for object users with read and write affinity to local clusters; and multi-site replication is a godsend for disaster recovery or archiving.
- ✔ **Cost-effectiveness:** Standard, commodity hardware allows you to tailor the optimal price/performance mix of standard servers and disk to each workload; thin provisioning enables you to create sparse block images in order to be able to over-provision the cluster. Support for heterogeneous hardware allows you to avoid having to replace older hardware as newer nodes are added; and erasure coding provides a cost-effective data durability option.

Chapter 2

Exploring Use Cases and Configuration Best Practices

In This Chapter

- ▶ Looking at real world OpenStack examples
 - ▶ Deep dive into hyperconverged and multi-region architectures
-

In this chapter, you discover some common OpenStack and Ceph use cases and configuration best practices.

Common Use Case Examples

Common use cases for the OpenStack architecture with Ceph storage technology include private and public cloud deployments, network functions virtualization (NFV) infrastructure, and application development environments.



To simplify recovery in a failure scenario so that a user only needs to re-bootstrap the application stack using Heat, configure the workload using the configuration management system and then start the application. When designing your cloud environment, you need to ensure the following:

- ✓ Images are a template of your applications.
- ✓ Application data is always hosted on Cinder block devices.

- ✓ Only ephemeral data is stored on the virtual machine root disk.
- ✓ Your application stack is managed by Heat orchestration.

On-premises private cloud

On-premises private cloud deployments are probably the most common OpenStack use case, requiring the cloud builder to expand her purview to include storage, specifically Ceph, in order to support an OpenStack infrastructure.

By adequately provisioning object and block volumes for virtual machines, software-defined, scale-out storage decreases the time it takes to roll out new services, such as sharing virtual machine images across cloud nodes with minimal downtime. This involves configuring the storage cluster and specifying the size of its disks for each new virtual machine within the OpenStack environment.



Red Hat OpenStack Platform provides simplified storage setup, initial configuration, and deployment with OpenStack Platform director. The director is a toolset for installing and managing a complete OpenStack environment.

Public cloud

Cloud service providers (CSPs) can deploy OpenStack with Ceph to build products and services similar to those offered by public cloud providers like Amazon, Google, and Microsoft. Ceph allows the administrator to offer object-, block- and file-storage-as-a-service with multi-tenancy security and charge-back capabilities.

NFV infrastructure

OpenStack is a critical element of an open NFV environment for telecommunications providers. NFV involves streamlining the network functions of telecom infrastructure by replacing legacy, proprietary hardware with elastic, virtualized environments that are multipurpose.

OpenStack with Ceph provides an OpenStack NFV solution with scale and performance — that’s affordable, on clustered, standard commodity server and disk hardware with high manageability and minimal downtime. Data is distributed among servers and disks dynamically, and the software relies on no single point of failure for maximum uptime.

Application development

OpenStack is used in AppDev environments that span development, deployment, and management, including environments involving containers where storage is once again essential. DevOps teams need the ability to choose the right type of storage for their needs and to be able to add and remove it programmatically.



In AppDev environments, Ceph provides two key benefits:

- ✔ Full API control and dynamic provisioning of any type of storage, be it object, block, or file. Developers can choose the right type of storage for their applications on an as-needed basis.
- ✔ The ability to provide persistent storage directly to applications running in containers.

Other possible AppDev use case might involve implementing Platform as a Service (PaaS) via Red Hat OpenShift Container Platform, Red Hat’s Docker-based Linux container solution with Kubernetes orchestration, on top of Red Hat OpenStack Platform. This combination helps developers develop, build, deploy, and manage containerized applications with full life cycle management.

Configuration Best Practices

Two common implementation scenarios for the cloud, NFV, and AppDev use cases include hyperconverged infrastructure and multi-region deployments.

Hyperconverged infrastructure

Hyperconverged infrastructures are becoming increasingly popular for cloud and NFV deployments due to their smaller initial deployment footprints, lower cost of entry, and maximized capacity utilization.



A server that runs both compute and storage processes is known as a *hyperconverged node*.

In a hyperconverged use case, Red Hat OpenStack Platform and Red Hat Ceph Storage are deployed such that the OpenStack Nova Compute Services and the Ceph Object Storage Daemon (OSD) services reside on the same node.

In order to achieve proper hyperconvergence, process resources must be correctly restrained and contained to an optimal level. Fortunately, OpenStack Nova has several capabilities that allow you to do so, including

- ✓ non-uniform memory access (NUMA) placement control
- ✓ CPU pinning and CPU ratio allocation
- ✓ kernel-based virtual machine (KVM) resource reservation

These capabilities allow you to ensure virtual machines are only using a subset of the bare metal server. You can then leave the rest of the resources to the Ceph OSD daemons.



An interesting approach to hyperconvergence is to use cgroup, a kernel feature that can limit and isolate resource consumption of OSDs. In this approach, cgroup makes sure that OpenStack Nova processes (virtual machines) and Ceph OSD processes do not conflict, thereby ensuring proper resource management and control.

Hyperconvergence can take several forms. The easiest (and default) form today is the one explained earlier in this section. However, as the trend toward containers continues to grow, another design has emerged: running both OpenStack Nova and Ceph OSD processes inside containers. Such a setup involves an OpenShift host, which is a lightweight PaaS, containing bare minimum packages and a container engine (Docker) that is optimized for running container workloads. Containers broaden hyperconvergence possibilities, including simplifying resource

isolation and restriction, enhancing upgrade capabilities, and leveraging an emerging industry standard to run applications.

A hyperconverged infrastructure deployment can start with as a few as six bare metal servers along with either a seventh bare metal server or a dedicated, separately hosted virtual machine for the deployment tools. These servers should be deployed in the following roles:

- ✔ One Red Hat OpenStack Platform director (can be virtualized for small deployments)
- ✔ Three cloud controllers which also host the Ceph Monitor services
- ✔ Three hyperconverged compute/Ceph Storage nodes

Six networks are recommended in the reference implementation and serve the following roles:

- ✔ **Ceph Cluster Network:** The Ceph OSDs use this network to balance data according to the replication policy. This private network only needs to be accessed by the OSDs.
- ✔ **Ceph Storage:** The Ceph OSDs public network used for Ceph OSDs public communications with the hypervisors. This network is also used by the Ceph monitor nodes.
- ✔ **External:** Red Hat OpenStack Platform director uses this network to download software updates for the overcloud, and the cloud operator uses this network to access Red Hat OpenStack Platform director to manage the overcloud. The controllers use this network to route traffic to the Internet for tenant services which are externally connected via reserved floating IP addresses, and overcloud users use this network to access the overcloud. The compute nodes don't need to be directly connected to this network because their instances will communicate via the tenant network to the controllers who will then route external traffic on their behalf to the external network.
- ✔ **OpenStack Internal API:** OpenStack provides both public facing and private API endpoints. This is an isolated network for the private endpoints.
- ✔ **OpenStack Tenant Network:** OpenStack tenants may create private networks implemented by VLAN (virtual local area network) or VXLAN (virtual extensible local area network) on this network.

✓ **Red Hat OpenStack Platform Director Provisioning:** OSP director serves DHCP (dynamic host configuration protocol) and PXE (pre-boot execution environment) services from this network to install the operating system and other software on the overcloud nodes, and this is the network that the cloud operator uses to access the overcloud nodes directly by SSH (secure shell) if necessary. The overcloud nodes, Controller/Mon, and Compute/OSD must be configured to PXE boot from this network so that they may be provisioned.

These six networks can be trunked as VLANs to connect to the servers as illustrated in Figure 2-1. The hyperconverged use case is currently limited as a reference architecture only.

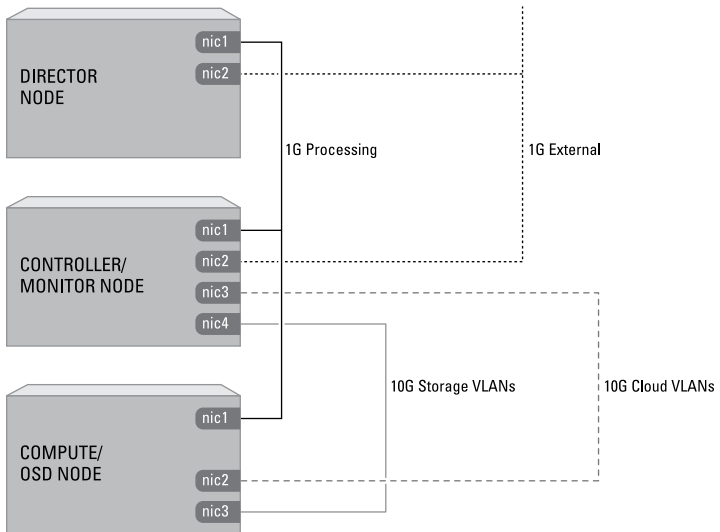


Figure 2-1: Network separation diagram.

Multi-region deployments

IT organizations require a disaster recovery strategy addressing outages with loss of storage or extended loss of availability at the primary site.

In this scenario, multiple isolated OpenStack environments are deployed independently with no services stretched between

them. Services at each site are maintained and kept highly available with Red Hat’s existing Pacemaker and HAProxy implementation, requiring no additional configuration.

The general idea is to seamlessly and transparently backup OpenStack images and block devices from one site to another. So in an event of a failure resources in site A can be manually brought online in site B. Client application load balancing should be done outside of OpenStack, as there is currently no OpenStack project offering a global load balancing mechanism. Instead, an external service, such as round robin DNS (domain name system), must be used to perform load balancing (see Figure 2-2).

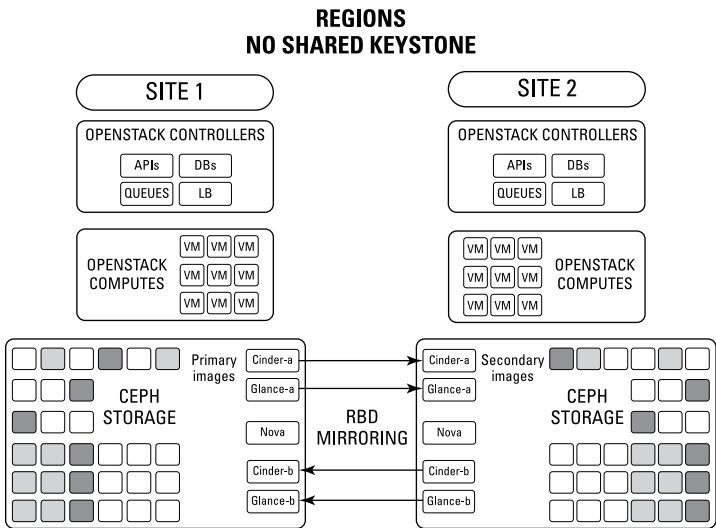


Figure 2-2: Multisite deployment with no shared keystone.

Sites can optionally be loosely coupled by aggregating the regions together by a shared-stretched Keystone (as well as Horizon if required) to create the perception of a unified global deployment (see Figure 2-3). In this topology, Keystone tokens need to be synchronized between sites, at least until fernet tokens are established as a stable feature. Users and tenants/projects will be able to consume resources from any region simultaneously and should do so for application resiliency.

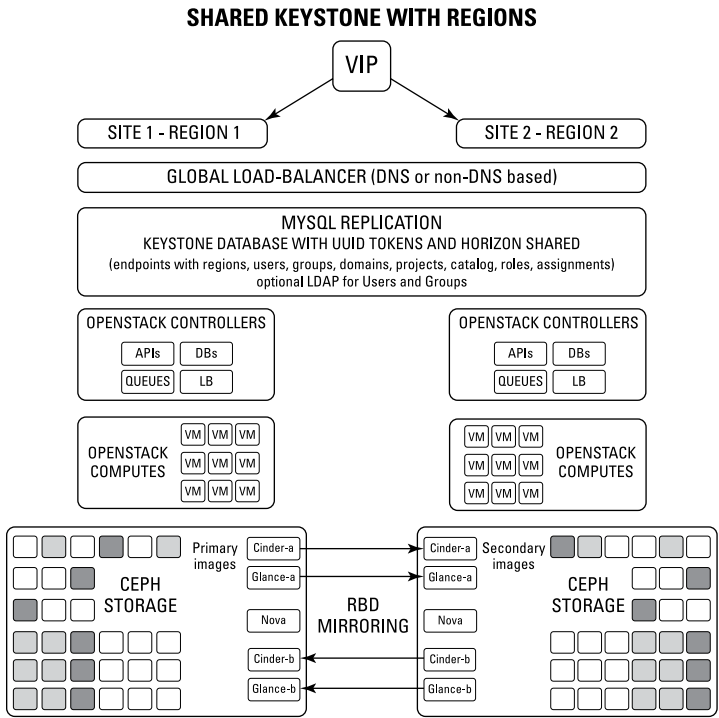


Figure 2-3: Multi-region deployment with shared keystone.



One benefit of using a shared keystone with OpenStack regions is the ability for the users to bootstrap their applications across three different locations using the same Keystone endpoint. Among other things, this brings more flexibility than having three different clouds operated by different individuals.

Chapter 3

Ten Reasons to Choose Red Hat Ceph Storage

In This Chapter

- ▶ Recognizing the business benefits
- ▶ Delivering scale-out enterprise performance
- ▶ Boosting security and stability

You know ‘em, you love ‘em. Here’s a list of ten reasons to choose Red Hat Ceph Storage for your OpenStack implementation, presented in classic *For Dummies* style:

- ✔ **Unified platform** that efficiently supports block, object, and file storage on standard servers and disks.
- ✔ **Tight “out-of-the-box” integration with OpenStack services**, including Nova, Cinder, Glance, Keystone, Manila, Ceilometer, and Swift. Ceph is the default storage backend when you deploy Red Hat OpenStack Platform.
- ✔ **User-driven storage** with 100 percent application programming interface (API) coverage that enables users to manage all cluster and object storage functions programmatically while gaining independence and speed by not having to manually provision storage.
- ✔ **Massive scale-out capability** supporting multiple petabytes of data, tiered I/O performance within a single cluster, and instantaneous booting of one to hundreds of virtual machines.
- ✔ **Maximum uptime** with no single point of failure, self-healing capabilities, and dynamic data distribution among servers and disks to reduce maintenance, and multisite replication for disaster recovery or archiving.



- ✔ **Robust ecosystem and no vendor lock-in** for long-term stability and a wide breadth of solutions.

As the primary contributor to the OpenStack and Ceph open source communities, Red Hat is one of many stakeholders dedicated to the ongoing innovation, pursuit of quality, and viability of these projects.

- ✔ **Enterprise readiness and peace of mind** with numerous consulting, services, training, and support options on stable code, backed by Red Hat Product Security, Red Hat Certification, and Red Hat Open Source and Quality Assurance Programs.
- ✔ **Reliable product life cycle management** with production-ready code, consistent packaging and distribution through Red Hat Satellite, and a well-defined release life cycle that's infrequent, hardened, and curated with extensive quality engineering and a committed three-year lifespan. Upgrades consist of timely, tested patches that are unlimited, not tied to a specific version or platform, and have a clearly defined, documented, and supported migration path.
- ✔ **Deployment resources and Red Hat subscription benefits** spanning reference architectures, performance and sizing guides, best practice guides, technical briefs, tutorials, real-time analytics, and troubleshooting devices to increase the chances of roll-out and production success.
- ✔ **Single source with added value** for open source solution components, all from Red Hat, including Red Hat Enterprise Linux, Red Hat Ceph Storage, Red Hat OpenStack Platform, Red Hat Cloud Infrastructure, and the newest integrated solution, Red Hat Cloud Suite. Every unique Red Hat OpenStack Platform account also receives — at no additional charge — 64 terabytes of Red Hat Ceph Storage with their subscription (resulting in 21 TBs of usable storage with the default recommended replica count of three.)

RED HAT®
CEPH STORAGE



**TRUSTED
FOR OPENSTACK**

Copyright © 2016 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and JBoss are trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.



redhat.

These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

Achieve maximum storage flexibility and performance with your OpenStack cloud

In this book, you'll learn about OpenStack® storage requirements and how Ceph addresses them with support for object, block, and shared file systems.

- **Learn how storage factors into a healthy OpenStack deployment** — examine the role of both ephemeral and persistent storage
- **Study real-world OpenStack use cases** — tap into private and public cloud, NFV, and application development environments
- **Explore configuration best practices** — delve into hyperconverged and multi-region implementation scenarios
- **Discover Red Hat Ceph Storage strengths** — leverage the enterprise-ready solution that's bundled with Red Hat OpenStack Platform



Open the book and find:

- How Ceph's extensible architecture integrates tightly with OpenStack
- How to deliver scale-out enterprise performance
- Ways to protect from failures using multisite replication for disaster recovery
- Ten reasons to choose Red Hat Ceph Storage

Go to [Dummies.com](https://www.dummies.com)
for videos, step-by-step examples,
how-to articles, or to shop!

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.