

**RED HAT
OPENSTACK PLATFORM**

**MAXIMIZE YOUR EVALUATION
WITH INSIGHT INTO TOP
FEATURES AND USES**

MAXIMIZE YOUR CLOUD EVALUATION

A guide to Red Hat OpenStack Platform features and use cases

This e-book provides an overview of select features, functionality, and use cases for Red Hat® OpenStack® Platform, answering *what?*, *why?*, and *how?* for each topic. Each section is independent, so you can use the table of contents to skip to sections that most interest you. Clicking on the cloud icon in the lower right corner of each page will return you to this table of contents.

Links to key documents and websites are included in each section, as well as in the general listing at the end of this e-book.

To get started:

- » Download the Red Hat OpenStack Platform evaluation at: access.redhat.com/products/red-hat-openstack-platform/evaluation.
 - » Review in-depth Red Hat OpenStack Platform product guides at: access.redhat.com/documentation/en-us/red_hat_openstack_platform.
-

Page 1

INTRODUCTION TO OPENSTACK

Page 2

WELCOME TO THE CLOUD

Page 3

RED HAT OPENSTACK PLATFORM SERVICES AND RELATIONSHIPS

Page 4

TOP FEATURES FOR INSTALLATION

Page 5

INTEGRATED DEPLOYMENT TOOLING

Page 7

TOP FEATURES FOR DAY-TO-DAY OPERATIONS

Page 8

AUTOMATED UPDATE AND UPGRADE PROCESSES

Page 9

INSTANCE LIVE MIGRATION

Page 10

UNIFIED LIFE-CYCLE MANAGEMENT

Page 13

COMPOSABLE ROLES AND SERVICES



MAXIMIZE YOUR CLOUD EVALUATION

A guide to Red Hat OpenStack Platform features and use cases

Page 14

**TOP FEATURES FOR NETWORK FUNCTIONS
VIRTUALIZATION (NFV)**

Page 15

VIRTUAL LOCAL AREA NETWORKS (VLANs)

Page 16

**HIGH-PERFORMANCE VIRTUAL
NETWORKING**

Page 17

**ENHANCED PACKET ROUTING AND
PROCESSING**

Page 18

SOFTWARE-DEFINED NETWORKING (SDN)

Page 19

**TOP FEATURES FOR BIG DATA AND
STORAGE**

Page 20

**APACHE HADOOP CLUSTER
IMPLEMENTATION**

Page 21

**PRODUCTION-GRADE, PROGRAMMABLE
STORAGE**

Page 23

**TOP FEATURES FOR SCIENCE AND
RESEARCH**

Page 24

PCI EXPRESS (PCIe) DEVICE ACCESS

Page 25

TOP FEATURES FOR DEVOPS

Page 26

**CONTAINER INTEGRATION AND
MANAGEMENT**

Page 27

GENERAL RESOURCES

Page 28

SCHEDULE A DISCOVERY SESSION



INTRODUCTION TO OPENSTACK

The top choice for cloud environments

OpenStack is an increasingly popular choice for cloud deployments across all industries and sizes of organizations. In fact, OpenStack is one of the fastest-growing open source communities in the world, with 83,767 community members and 672 supporting companies in 189 countries contributing more than 20 million lines of code.¹

When building a private or public Infrastructure-as-a-Service (IaaS) cloud environment, organizations choose OpenStack to:²

- » Increase operational efficiency.
- » Accelerate their ability to innovate.
- » Avoid vendor lock-in.
- » Save money.
- » Standardize on the same open platform and application programming interfaces (APIs) that power a global network of public and private clouds.
- » Achieve security and privacy goals.
- » Attract top technical talent.

OPENSTACK BY THE NUMBERS¹

20 MILLION
{{{ lines of code }}}



83,767
members

672
companies



189
countries

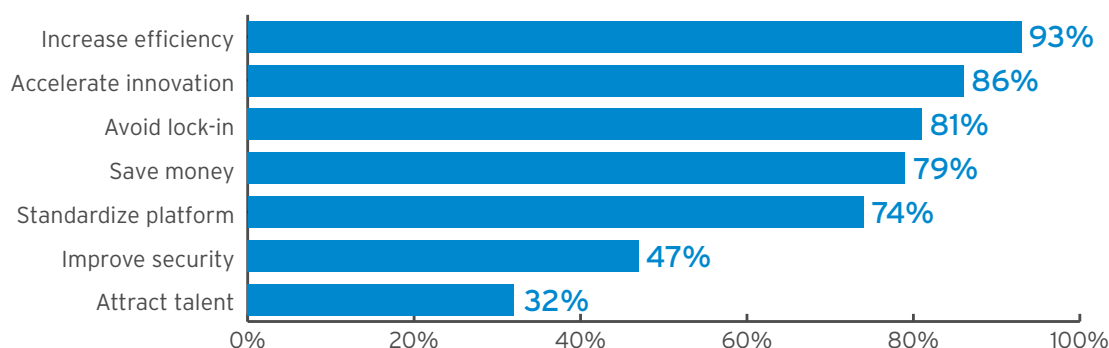


Figure 1. Top reasons for choosing OpenStack²

¹ OpenStack Foundation. From openstack.org as of December 17, 2017.

² "OpenStack User Survey." The OpenStack Foundation. November 2017. openstack.org/assets/survey/OpenStack-User-Survey-Nov17.pdf.

WELCOME TO THE CLOUD

The power and stability of Red Hat OpenStack Platform

As business competition and disruption increase across industries, IT organizations are challenged to respond faster to new, growing, and fast-changing customer demands. An OpenStack cloud environment can provide agility and speed, but not all OpenStack distributions can meet the additional demands of production-grade deployment, performance, scalability, and security. Red Hat OpenStack Platform can.

Red Hat OpenStack Platform combines the power of Red Hat Enterprise Linux® with OpenStack technology developed by Red Hat to deliver a foundation for building private or public IaaS cloud environments. It offers a highly scalable, fault-tolerant platform for running cloud-enabled workloads. Red Hat OpenStack Platform also lets you dynamically expand capacity across cloud environments to meet demand peaks without downtime or complex processes.


OpenStack depends on Linux as the operating environment for its services, to access hardware resources, and for third-party integration with existing or new systems. OpenStack also depends on Linux for system-wide performance, scalability, and security and as an operating system for the guest applications running in these virtual environments. This is why Red Hat—recognizing the critical interdependencies between the two—co-engineered Red Hat OpenStack Platform and Red Hat Enterprise Linux together with closely aligned product engineering teams. Paired with Red Hat CloudForms' open, hybrid management for workloads and infrastructure, as well as Red Hat Ceph Storage's highly redundant, software-defined storage, Red Hat OpenStack Platform is an integrated, optimized, and managed foundation for production-ready cloud environments.

Red Hat OpenStack Platform is also integrated with other products in Red Hat's software stack. Red Hat Ansible® Automation—simple, flexible IT automation technology—works with Red Hat OpenStack Platform director to simplify deployment, upgrades, and overall management. Red Hat OpenStack Platform also makes an ideal foundation for Platform-as-a-Service (PaaS) deployment with Red Hat OpenShift Container Platform—a platform for building, deploying, and managing container-based applications consistently across hybrid cloud and multicloud infrastructures.

RED HAT® OPENSTACK® PLATFORM

KEY FEATURES

- » Co-engineered performance
- » Built-in security
- » Unified cloud management
- » Scale-out services, including compute, network, and storage
- » Freedom from vendor dependency
- » Production-friendly life-cycle and support options



OPENSTACK SERVICES AND RELATIONSHIPS

A modular, interoperable architecture for flexibility

Red Hat OpenStack Platform features a modular architecture based on a variety of interoperable services for compute, networking, storage, management, orchestration, and more. These services can be used in any combination to align your cloud infrastructure with the needs of your organization.

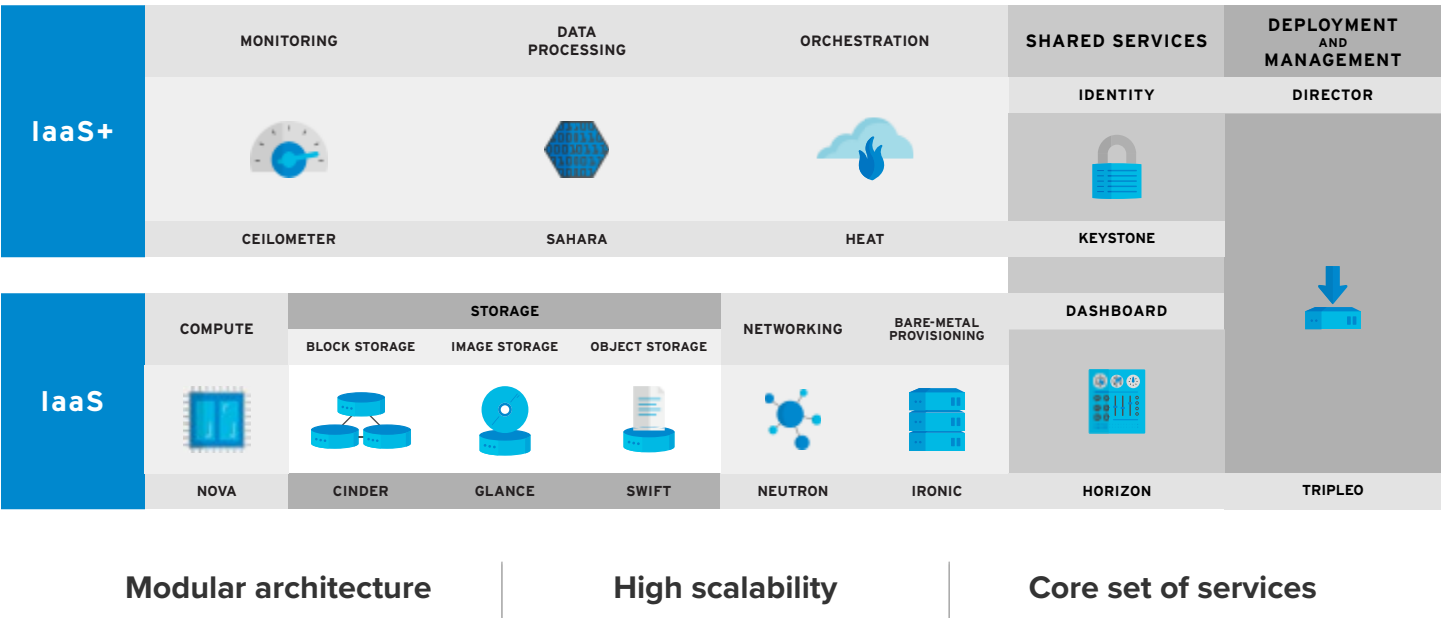


Figure 2. Red Hat OpenStack Platform modular architecture with interoperable services



TOP FEATURES FOR INSTALLATION

As the first step in using any technology, installation is critical. Ensuring components are installed and configured correctly can be the difference between a successful technology project and a failed one. With a wide variety of services and options, installing OpenStack can prove daunting. Red Hat OpenStack Platform director alleviates installation and ongoing operational challenges through deployment and life-cycle management automation.



INTEGRATED DEPLOYMENT TOOLING

Red Hat OpenStack Platform director

OpenStack cloud deployments are complex—there is no standard cloud configuration that fits all use cases. Differing requirements, expectations, hardware, networking, storage, and datacenter services make it impossible to define a single, standard deployment. In addition, OpenStack deployments themselves are complicated. According to the OpenStack Foundation, the average OpenStack deployment deploys 11 OpenStack services.³

Director is the integrated deployment and life-cycle management tool included with Red Hat OpenStack Platform. Integration with Red Hat's management and automation tools allows you to use an infrastructure-as-code approach to cloud management. Use director with Red Hat Ansible Automation to automate your cloud infrastructure, Red Hat Satellite to manage subscriptions and content, and Red Hat CloudForms to perform advanced Day 2 operations. Additionally, based on the upstream community's TripleO project, director uses OpenStack services to install and deploy your Red Hat OpenStack Platform environment. As a result, installation and management are aligned with and benefit from improvements to the overall OpenStack platform.

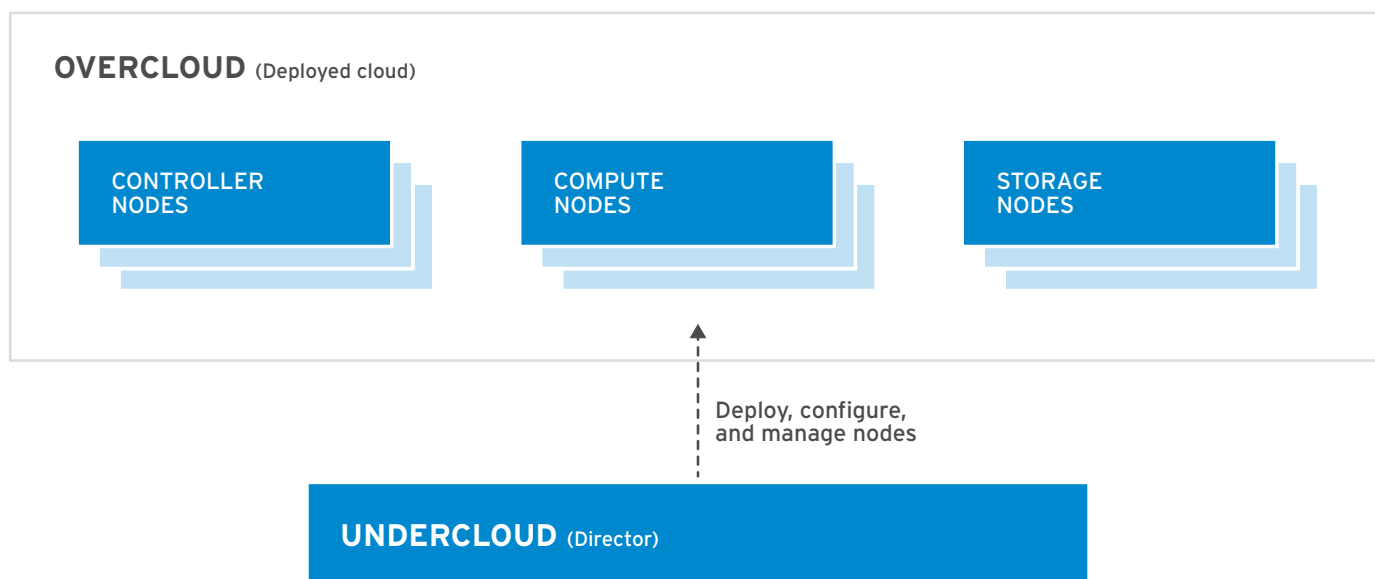


Figure 3. Red Hat OpenStack Platform director installs, configures, and manages OpenStack nodes

³ "OpenStack User Survey." The OpenStack Foundation. November 2017. openstack.org/assets/survey/OpenStack-User-Survey-Nov17.pdf.

INTEGRATED DEPLOYMENT TOOLING

Red Hat OpenStack Platform director

Using repeatable processes, Red Hat OpenStack Platform director can orchestrate deployment and setup of the majority of cloud configurations and use cases. Input specific parameters into predefined configuration and template files. You can also customize your environment to cover corner use cases using a flexible template framework that can be automated and plugged into a continuous integration/continuous deployment (CI/CD) system. Red Hat OpenStack Platform director will provision hardware with an operating system, configure and enable OpenStack components, and make services available for use. During the process, preflight checks validate the parameters and inflight checks ensure components are configured as requested. Post-deployment checks ensure that your environment is functioning according to the requested configurations and the OpenStack specification. You can also choose to benchmark the new configuration to identify potential performance bottlenecks using the Tempest and Rally services. A similar process applies to configuration of minor updates, major upgrades, and changes to your Red Hat OpenStack Platform environment.

After installation, Red Hat OpenStack Platform director gives you a simple, streamlined way to interact with your environment. Access controls ensure only authorized users can make changes. Director also acts as an abstraction layer, preventing users from directly accessing OpenStack servers. Users interact with Red Hat OpenStack Platform via your overcloud and Horizon, while administrators and operations interact with OpenStack Platform via director.

WHAT?

Red Hat OpenStack Platform director—the integrated deployment and life-cycle management tool for Red Hat OpenStack Platform—automates and streamlines cloud deployment and administration.

WHY?

OpenStack clouds can be very complex and ensuring proper installation can be challenging—especially if you don't have extensive OpenStack expertise.

HOW?

Director integrates with Red Hat Ansible Automation and includes a comprehensive set of Ansible Playbooks for installation, deployment, and configuration operations.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for these guides:

- » Director installation and usage
- » Introduction to the OpenStack dashboard
- » Advanced overcloud customization





TOP FEATURES FOR DAY-TO-DAY OPERATIONS

Once your cloud environment is deployed, you need to use it to support your business. Day-to-day operations become essential for delivering services and resources, ensuring security, and maintaining compliance and governance standards. Integration with Red Hat Ansible Automation and Red Hat CloudForms, as well as live migration capabilities, help you operate your Red Hat OpenStack Platform cloud environment efficiently and effectively.



AUTOMATED UPDATE AND UPGRADE PROCESSES

Red Hat OpenStack Platform director and Red Hat Ansible Automation

Changes, updates, and upgrades are a critical part of infrastructure maintenance and can be particularly challenging in complex production environments. As business and IT requirements change, you need to adjust infrastructure and service configurations, add new services, and scale resources. In addition, new OpenStack versions are released on a six-month cadence. Staying current with the latest patches and platform versions helps ensure that your operating environment continues to support your business with innovation, efficiency, and protection. Even so, these modifications often increase production risks, as new configurations, software features, and platform versions must be proven out and configured correctly. Flawed adjustments can result in downtime, inefficiencies, security vulnerabilities, and service delays.

Manual processes can exacerbate these risks, as the chance of inconsistencies and errors increases. Automated technologies can help reduce risks associated with configuration changes, updates, and upgrades.

Red Hat Ansible Automation supplies a simple, powerful, and agentless framework for network, server, and application automation. With Ansible Automation, you can add stability, scalability, and repeatability to your cloud management processes. Integration between Red Hat OpenStack Platform director, Ansible Automation, and all OpenStack services lets you automate updates, upgrades, and other Day 2 operations to increase consistency, speed, and understanding. Some of the tasks that director simplifies and automates are:

- » Update patch installation.
- » Major version upgrades.
- » Environment scale-out.
- » Post-deployment configuration changes.
- » Monitoring, through both Red Hat CloudForms and Red Hat Ansible Automation.
- » Backups.

WHAT?

Integration between Red Hat OpenStack Platform director and Red Hat Ansible Automation automates, simplifies, and speeds updates, upgrades, and day-to-day cloud management operations.

WHY?

Manual processes can be slow and error-prone. Automation ensures better consistency, accuracy, and speed for updates, upgrades, and day-to-day operations.

HOW?

Director integrates with Red Hat Ansible Automation and all OpenStack services. Red Hat OpenStack Platform includes a comprehensive set of Ansible Playbooks for updates, upgrades, and other Day 2 operations.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for these guides:

- » [Upgrading Red Hat OpenStack Platform](#)

For information about life-cycle and update policies, see [Red Hat OpenStack Platform life cycle](#).



INSTANCE LIVE MIGRATION

Red Hat OpenStack Platform director and nova live-migration

Over time, you'll inevitably need to perform maintenance on your cloud environment. Ideally, you'll need to do this without taking virtual machines (VMs) and workloads down so that you can continue to support your users and business at all times. Typically, this involves manually moving each VM to another host, which can be both time-consuming and error-prone.

Instance live migration capabilities, included in Red Hat OpenStack Platform, let you seamlessly move VMs from one host to another without downtime. Using the [nova live-migration](#) command, you can automatically move all VMs on a given host to other hosts in the cluster, saving you time and ensuring continuity. This allows you to more easily perform host and hypervisor maintenance. To take advantage of instance live migration, you'll need to set up shared storage between hosts, as storage that is attached only to a particular host will be inaccessible when that host is unavailable.

Instance live migration also plays a role in use cases that require instance high-availability and in supporting stateful applications. If a hypervisor fails, all VMs on that host can be evacuated. You can configure this through Red Hat OpenStack Platform director's command line interface (CLI) or through a graphical user interface (GUI) in Red Hat CloudForms, which is included in your Red Hat OpenStack Platform subscription.

WHAT?

Instance live migration capabilities let you move a running VM from one host to another without downtime.

WHY?

Your environment will need periodic maintenance, but you still need to provide services to users at all times.

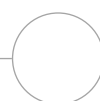
HOW?

Shared storage is required. The host evacuate command moves all VMs on a failed host to other hosts in the cluster without manual intervention.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for this guide:

» [Instances and images guide](#)



UNIFIED LIFE-CYCLE MANAGEMENT

Red Hat CloudForms

Resource and infrastructure management is critical for supporting business operations. Workloads and services must be deployed, monitored, moved, and retired throughout their life cycle. Inefficient processes and multiple, disparate management tools can delay service and resource provisioning, decrease resource utilization and productivity, and result in inconsistencies across your environment. Unifying resource management under a single tool can increase efficiency, visibility, and insight across infrastructures.

Red Hat CloudForms is a cloud management platform (CMP) that provides unified, comprehensive, consistent management across multivendor virtualized, private, hybrid, and public cloud, and container-based environments. Each Red Hat OpenStack Platform subscription includes a basic subscription for Red Hat CloudForms that allows you to manage your Red Hat OpenStack Platform environment. A full Red Hat CloudForms subscription can be added to extend and unify management across all parts of your hybrid and multi-cloud environment.

Through both undercloud and overcloud integration, Red Hat CloudForms gives you simple, complete life-cycle management of your Red Hat OpenStack Platform environment and resources. It provides self-service capabilities through automated provisioning and policy enforcement for accelerated service delivery. It also gives IT managers complete operational and life-cycle control over the deployed services. As a result, you can improve service levels while retaining control, security, and compliance. Red Hat CloudForms takes an industry-first, Red Hat Ansible Automation-based approach to multi-cloud management, natively supporting Ansible Playbooks to make it easier to deploy IT services across clouds. This approach allows you to automate a wide range of policies and processes with minimal coding or scripting. Finally, continuous monitoring and automation capabilities let you trigger events based on collected data—without manual intervention.

WHAT?

Red Hat CloudForms provides simplified, unified life-cycle management of OpenStack resources and Day 2 operations.

WHY?

Manually performing day-to-day operations in a complex OpenStack environment can result in delays, security risks, and compliance issues—negating many of the benefits of a cloud environment.

HOW?

Red Hat CloudForms runs as an external appliance with *providers* connected to the overcloud and undercloud. Storage and network provider connections are created automatically.

LEARN MORE

Visit [product documentation for Red Hat CloudForms](#) and look for this guide:

- » Installing Red Hat CloudForms on Red Hat OpenStack Platform



UNIFIED LIFE-CYCLE MANAGEMENT

Addressing OpenStack management challenges

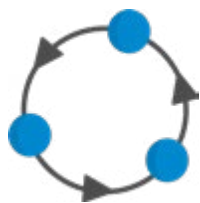
Red Hat CloudForms addresses many of the top challenges of managing OpenStack environments, including operational efficiency, service health, security and compliance, and financial management.



OPERATIONAL EFFICIENCY

Gain visibility into and control over your OpenStack environment.

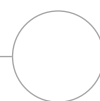
- » Continuous discovery and inventory of infrastructure resources in near real time gives you visibility into your deployment.
- » Resource relationship discovery and topology visualization helps you understand where resources are located and how they are connected.
- » Compute, storage, and network monitoring and data analysis help you troubleshoot issues across your stack.
- » Automation capabilities through control policies and user interface extensions let you streamline common tasks.



SERVICE HEALTH

Control configuration drift and ensure services have the resources they need to meet performance requirements.

- » Automatic configuration tracking and enforcement curbs drift.
- » Performance and resource event alerting lets you know of issues faster.
- » Right-sizing recommendations helps you optimize resource utilization.
- » Optimization and planning capabilities let you track and project capacity needs.



UNIFIED LIFE-CYCLE MANAGEMENT

Addressing OpenStack management challenges



SECURITY AND COMPLIANCE

Protect your OpenStack environment from threats and ensure compliance with enterprise policies and regulations.

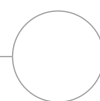
- » Continuous monitoring and flagging of non-compliant VMs ensure resources comply with enterprise policies.
- » Package, init process, and file retrieval lets you know if users are running untrusted software.
- » User and group data collection helps you discover unauthorized users.
- » Configuration tracking gives you visibility into changes over time.
- » Granular role-based access control (RBAC) and tenant synchronization let you manage user permissions.



FINANCIAL MANAGEMENT

Better understand and control costs within your OpenStack environment.

- » Tracking of top CPU, memory, and cost users shows you who is consuming the most resources.
- » Financial modeling and analysis help you understand your costs.
- » Chargeback and showback capabilities—by VM or tenant—let you pass costs through to users based on utilization.



COMPOSABLE ROLES AND SERVICES

Independently scalable services

Though cloud environments boast massive scalability, older OpenStack distributions used a controller architecture that limited scalability to around three controller nodes and 300 compute nodes for guaranteed performance. API, database, and message bus services were combined into a single, monolithic controller node. Scaling beyond three of these nodes delivered diminishing scalability.

Red Hat OpenStack Platform 10 and later versions feature composable roles and services. This architecture separates the monolithic controller node into its core services, allowing them to scale independently and eliminating many performance issues for large-scale OpenStack deployments. Composable roles give you more control over how OpenStack services are allocated within your environment. You can split services into individual roles or combine them into new types of services and nodes. For example, you might configure a networker node using the Neutron and Designate projects, or a hyperconverged compute and storage node using Nova, Glance, Swift, and Cinder. As a result, you can distribute services in a completely flexible manner and scale services—including the control plane—in any direction.

WHAT?

Composable roles deliver more control over the distribution of OpenStack services within your environment.

WHY?

Scale your environment in any direction without compromising performance.

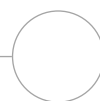
HOW?

Create, deploy, and scale custom OpenStack roles from individual services or combinations of services.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for this guide:

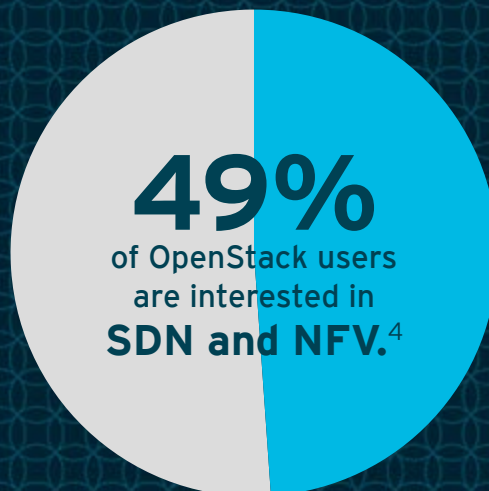
- » Advanced overcloud customization: Composable services and custom roles





TOP FEATURES FOR NETWORK FUNCTIONS VIRTUALIZATION (NFV)

NFV and SDN are key initiatives for the fast-changing telecommunications industry. Both technologies can help communications service providers (CSPs) deliver new services faster and operate more efficiently and flexibly. NFV implements network functions as VMs on a highly scalable, cloud-based infrastructure, rather than as dedicated physical hardware with proprietary components. SDN creates an agile, programmable network for service delivery and control. OpenStack plays a central role in deploying NFV and SDN and 49% of OpenStack users say that they are interested in SDN and NFV.⁴



⁴ "OpenStack User Survey." The OpenStack Foundation. April 2017. openstack.org/assets/survey/April2017SurveyReport.pdf.



VIRTUAL LOCAL AREA NETWORKS (VLANs)

VLAN-aware VMs and VLAN trunk ports

Many NFV use cases—such as virtual customer premise equipment (vCPE)—require network connections to a large number of devices. Even with VLANs, managing direct connections for each device is challenging at best. Running VLAN-aware VMs in Red Hat OpenStack Platform can alleviate manageability issues through VLAN trunking.

VLAN trunking eliminates the need to provision a specific virtual network interface card (vNIC) for each VLAN, resulting in fewer VLANs per VM and simpler NFV deployments. VLAN trunk ports combine and encapsulate the multiple VLAN connections in each OpenStack instance and direct all traffic from the instance over a single vNIC. Once a packet reaches the main network fabric, the encapsulation is remapped and redirected to the appropriate endpoint. In an NFV environment, most virtual network functions (VNFs) expect this type of network behavior as it emulates physical networking environments.

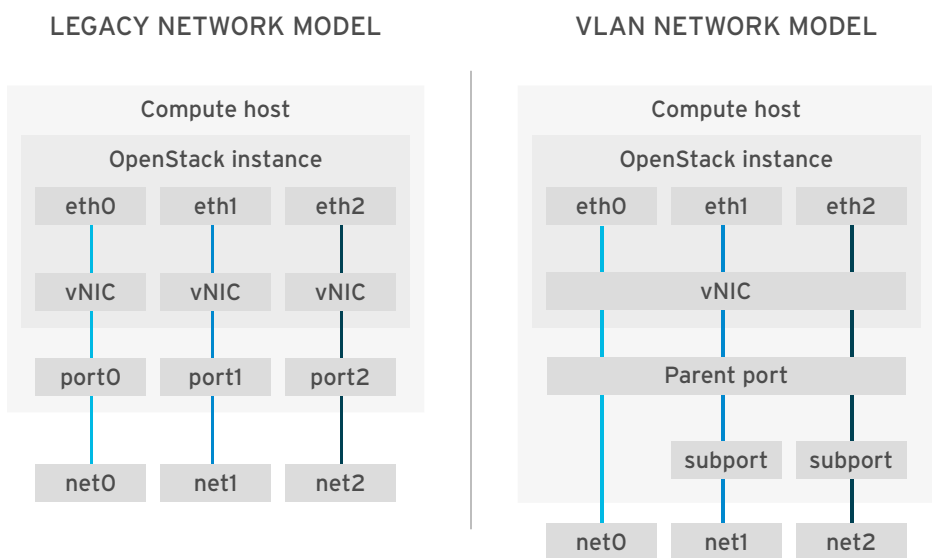


Figure 4. Legacy versus VLAN network models

WHAT?

VLAN trunk ports let you connect multiple VLANs to a single vNIC on an instance.

WHY?

Networks with many device connections are complex and difficult to manage. VLAN trunking emulates a true networking environment and simplifies deployments.

HOW?

VLAN trunk ports combine and encapsulate the multiple VLAN connections in each OpenStack instance and direct all traffic from the instance over a single vNIC. Once a packet reaches the main network fabric, the encapsulation is remapped and redirected to the appropriate endpoint.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for this guide:

» Networking guide: VLAN-aware instances

To learn more about setting up trunk ports, see [How to create VLAN-aware VMs/neutron trunk ports in Red Hat OpenStack Platform](#).

HIGH-PERFORMANCE VIRTUAL NETWORKING

Single root input/output virtualization (SR-IOV)

High performance is a critical requirement for NFV environments, and your network configuration can have a large impact on VM performance. Red Hat OpenStack Platform supports SR-IOV to deliver higher performance for VMs.

SR-IOV evolved from PCI pass-through concepts, but instead of simply dedicating a physical network interface card (NIC) to a particular VM, SR-IOV allows a physical NIC to be shared across multiple VMs. Through the manufacturer's driver, you can virtualize a NIC to create multiple virtual functions on top of the physical function. These virtual functions are then attached to VMs using Red Hat OpenStack Platform director. This technique bypasses the Open vSwitch (OVS) and hypervisor layers, reducing the performance impact of virtualization and delivering better performance for high input/output operations per second (IOPS) applications.

WHAT?

SR-IOV allows a physical NIC to be shared between multiple VMs.

WHY?

By bypassing the OVS and hypervisor layers, SR-IOV provides fast network access for workloads with high IOPS requirements.

HOW?

Set up virtual functions through your NIC's driver, then attach VMs using Red Hat OpenStack Platform director.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for these guides:

- » Network functions virtualization planning and prerequisites guide: Planning your SR-IOV deployment
- » Network functions virtualization configuration guide: Configure SR-IOV support for virtual networking



NOTE:

The maximum number of virtual functions is determined by the NIC driver.



ENHANCED PACKET ROUTING AND PROCESSING

Open vSwitch (OVS) and Data Plane Development Kit (DPDK)

OVS and DPDK are essential components for high-performance NFV environments. OVS performs virtual switching throughout the environment. The DPDK library allows VNF software to route network packets around the Linux operating system kernel, which can significantly improve packet processing performance. Combining OVS with DPDK allows OVS to operate in a poll-based mode—rather than the default interrupt-based mode—reducing CPU utilization by kernel processes and improving overall latency and performance. In contrast to SR-IOV, this combination also enables live migration of VMs, giving you more flexibility with nearly equal performance.

Red Hat OpenStack Platform lets you easily enable OVS with DPDK across your environment. A DPDK-enabled version of OVS is supplied with your download and can be configured using director during installation.

WHAT?

Together, OVS and DPDK improve packet routing and processing for OpenStack-based NFV environments.

WHY?

Performance is critical for NFV environments and the standard OVS implementation can be inefficient.

HOW?

When combined with OVS, DPDK routes network packets around the Linux operating system kernel, minimizing overhead and allowing OVS to operate in a more efficient, poll-based mode.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for these guides:

- » Network functions virtualization planning and prerequisites guide: planning your OVS-DPDK deployment
- » Network functions virtualization configuration guide: configure OVS-DPDK support for virtual networking



NOTE:

While DPDK is supported by the open source community, it currently provides the best performance on Intel processors.

SOFTWARE-DEFINED NETWORKING (SDN)

SDN controller certifications and integrations

Efficient networking is essential to getting the most from your cloud infrastructure. Complex, static networking models require tedious, error-prone manual programming, which increases operational costs (OpEx), reduces efficiency, and delays the availability of network services for cloud applications. SDN can alleviate these challenges. Adding SDN capabilities to your cloud environment can make your network as flexible and responsive as your virtualized server and storage infrastructure, so you can improve productivity and more easily align your cloud infrastructure with ever-changing requirements.

While the standard Red Hat OpenStack Platform networking components include everything you need to provide basic networking capabilities for VMs and instances, they do not integrate with existing networking infrastructure—your OpenStack environment is a networking island. For organizations that wish to integrate their OpenStack environment into their existing network, Red Hat validates and certifies stable, reliable Red Hat OpenStack Platform operation with industry-leading SDN controllers, including OpenDaylight. This allows you to fully connect your OpenStack environment with your SDN controller, giving it visibility into and control over OpenStack networking, utilization, and policies. Integration with your switching environment creates a single, consistent fabric across your datacenter.

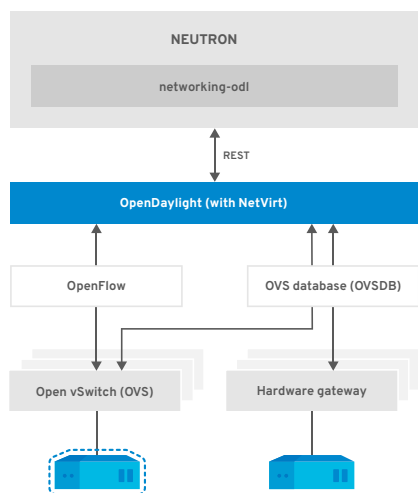


Figure 5. Red Hat OpenStack Platform and OpenDaylight cooperation

WHAT?

Integrate your Red Hat OpenStack Platform environment into your existing software-defined network.

WHY?

SDN integration gives you programmatic control over your OpenStack environment, allowing you to gain visibility, automate management, and create a single, consistent network fabric.

HOW?

Red Hat tests and certifies Red Hat OpenStack Platform with leading SDN controllers. Use director to enable SDN integration during deployment.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for these guides:

- » Networking guide
- » Red Hat OpenDaylight product guide
- » Red Hat OpenDaylight installation and configuration guide



TOP FEATURES FOR BIG DATA AND STORAGE

Data is an increasingly valuable asset for organizations, but, to be useful, you must be able to readily analyze and derive insight from it. This requires real-time access to massive amounts of data, as well as advanced big data analytics tools. To achieve success, you need a foundation that integrates multiple data sources and can transition workloads across on-premise and cloud boundaries.

WHAT IS BIG DATA?

Big data is popularly defined by these characteristics, known as the three Vs: Massive volumes of data in a variety of nonstandard formats that are processed at a high velocity.



VOLUME



VARIETY



VELOCITY



By 2025, IDC estimates that

180 ZETTABYTES

of data will be created each year.⁵

⁵ "IoT Mid-Year Update From IDC And Other Research Firms." *Forbes*. August 5, 2016.
forbes.com/sites/gilpress/2016/08/05/iot-mid-year-update-from-idc-and-other-research-firms/#531e061655c5.



APACHE HADOOP CLUSTER IMPLEMENTATION

OpenStack Data Processing (Sahara project)

Red Hat OpenStack Platform provides a scalable, flexible, high-performance foundation for big data workloads, but big data requires specific analytics frameworks to deliver results. The OpenStack Data Processing service—known as the Sahara OpenStack project—provides a simple means to provision and scale data processing frameworks—like Apache Hadoop, Spark, and Storm—within your OpenStack environment. Data Processing is fully integrated into the OpenStack ecosystem, so you can easily deploy, manage, and submit jobs to data processing clusters using the OpenStack dashboard. Simply specify configuration parameters, including framework version, cluster topology, and node hardware details.

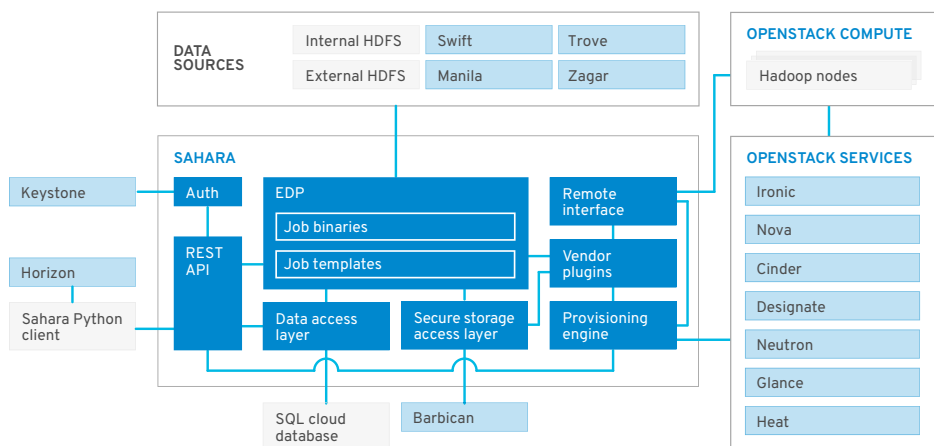


Figure 6. Data processing architecture within OpenStack

WHAT?

OpenStack Data Processing lets you easily provision and scale Apache Hadoop clusters to process large datasets like Cloudera, Hortonworks, and MapR.

WHY?

Big data analysis requires fast access to large datasets and specific processing frameworks.

HOW?

Manage the entire Hadoop data processing workflow—from configuration to launching and running jobs—through the OpenStack dashboard.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for this guide:

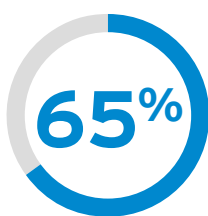
» [OpenStack Data Processing](#)

PRODUCTION-GRADE, PROGRAMMABLE STORAGE

Red Hat Ceph Storage

Real-time access to data is critical for big data analysis and insight. As a result, you need a storage platform that can scale dynamically and massively, provide fast, consistent access to data, and respond quickly to spikes in demand. Red Hat Ceph Storage is an open, highly scalable, programmable storage solution for modern workloads like cloud infrastructure, data analytics, media repositories, and backup and restore systems. Consistent with OpenStack deployment requirements, this software-defined platform provides a scale-out architecture for natural and containerized deployment of block, object, and file storage on industry-standard hardware.

Integration between Red Hat OpenStack Platform and Red Hat Ceph Storage allows you to instantaneously boot one to hundreds of VMs, which are made readily accessible to cloud users. This also simplifies and streamlines installation, upgrades, and updates. Using Red Hat OpenStack Platform director, you can deploy, manage, and upgrade Red Hat Ceph Storage and compute resources in a hyperconverged deployment. Additionally, seamless coordination between Red Hat Ceph Storage and the OpenStack Glance, Cinder, and Nova services enables rapid backups.



of OpenStack deployments
use Ceph storage.⁶

2/3

of organizations cannot cope
with workloads larger than
PETABYTE-SCALE.⁷

WHAT?

Red Hat Ceph Storage provides highly scalable, software-defined storage for Red Hat OpenStack Platform deployments.

WHY?

Storage performance and flexibility are critical for data access, processing, and analysis.

HOW?

Integration between Red Hat Ceph Storage and Red Hat OpenStack Platform lets you deploy, scale, and manage your storage back end just like your cloud infrastructure. Red Hat Ceph Storage provides reliable, flexible storage for key OpenStack services.

LEARN MORE

Visit [product documentation for Red Hat OpenStack Platform](#) and look for these guides:

- » Storage guide
- » Integrating an overcloud with an existing Red Hat Ceph cluster
- » Deploying an overcloud with containerized Red Hat Ceph Storage
- » CephFS back end guide for the shared file system service


⁶ "OpenStack User Survey." The OpenStack Foundation. April 2017.
openstack.org/assets/survey/April2017SurveyReport.pdf

⁷ "The storage conundrum." Vanson Bourne for Red Hat. 2017.
redhat.com/en/technologies/storage/vansonbourne.

RED HAT CEPH STORAGE

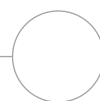
Top features for scalable OpenStack storage

Red Hat Ceph Storage includes many key features for Red Hat OpenStack Platform deployments.



- » Support for **block, object, and file storage** in one unified platform
- » Alignment with OpenStack's **modular architecture and services** like Nova, Cinder, Manila, Glance, Keystone, Ceilometer, and Swift
- » Management of **petabytes of data** with little to no performance degradation
- » User-driven storage life-cycle management with **100% API coverage**
- » **Common platform for workloads and applications** through Network File System (NFS), Internet Small Computer Systems Interface (iSCSI), and object support
- » **Security, reliability, and availability** capabilities, including geo-replication and disaster recovery
- » Cost-effective, high **performance**
- » **Self-healing and adaptation** to changing application and deployment demands
- » Shared **open development process** and large partner ecosystem

Learn about OpenStack and Ceph configuration best practices in [OpenStack Storage for Dummies](#).





TOP FEATURES FOR SCIENCE AND RESEARCH

Science and research organizations continually push the limits of compute infrastructure in their quest for knowledge. Historically, these organizations used high-performance computing (HPC) techniques and monolithic supercomputers to run data- and compute-intensive workloads. However, even with advanced processors and networking, the powerful servers needed for these workloads can be costly.

Cloud computing provides a modern, powerful, and cost-effective infrastructure for running researchers' diverse workloads. Dynamic scalability improves resource utilization and enables economical infrastructure expansion. Rapid resource provisioning and adaptable performance deliver faster results and iterations. New innovation can be adopted quickly to further advance research.

The modular architecture of OpenStack is ideal for varied science and research workloads and 10% of respondents to the November 2017 OpenStack Foundation user survey are using OpenStack for academic and research purposes.⁸

10% of OpenStack users employ it for
academic and research purposes.⁸

⁸ "OpenStack User Survey." The OpenStack Foundation. November 2017. openstack.org/assets/survey/OpenStack-User-Survey-Nov17.pdf.



PCI EXPRESS (PCIe) DEVICE ACCESS

Direct pass-through capabilities

Specialized science and research workloads often need access to additional hardware resources—like graphics processing units (GPUs) and InfiniBand interconnects—to access, process, and visualize large amounts of data and results. These resources are typically attached through PCIe interconnects. Red Hat OpenStack Platform provides support for these PCIe devices through direct pass-through capabilities.

Using Red Hat OpenStack Platform director, you can configure OpenStack compute instances to directly access PCIe devices. Fast datapath devices and information are automatically gathered and stored. Director uses this information, along with simple inputs from cloud administrators, to automatically calculate ideal virtualization settings at deployment.

WHAT?

Direct pass-through capabilities allow OpenStack instances to natively access PCIe devices like GPUs and network interconnects.

WHY?

Many science and research application require access to additional hardware resources attached through PCIe interconnects.

HOW?

Red Hat OpenStack Platform director lets you configure OpenStack instances with direct pass-through access to specific resources and then automatically provisions those resources.

LEARN MORE

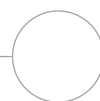
Visit [product documentation for Red Hat OpenStack Platform](#) and look for this guide:

» Bare metal provisioning

Additional information may be found in the Red Hat Knowledgebase and upstream documentation:

» [How to deploy and use nova PCI passthrough with Red Hat OpenStack Platform 10](#)

» [Enable PCI passthrough \(compute\)](#)





TOP FEATURES FOR DEVOPS

Software is at the core of business today. DevOps is an approach to culture, automation, and platform design to provide better business value and responsiveness in a software-based world. This approach fosters collaboration and communication between all stakeholders in the development and deployment process. It brings groups across the organization—specifically development and operations—together to increase the speed and flexibility with which new features and services are delivered. As a result, you can accelerate time to market for new offerings, improve productivity and efficiency, and reduce overall application costs.



CONTAINER INTEGRATION AND MANAGEMENT

Red Hat OpenShift Container Platform

According to OpenStack.org's April 2017 user survey report, containers are the most interesting emerging technology for OpenStack users, with 75% of users saying that they are interested in containers.⁹ Red Hat OpenStack Platform provides full container support—in fact, OpenStack services actually run in containers. This simplifies roll-out and distribution of services across nodes, as well as updates and upgrades, and makes it easier to integrate container technologies into your cloud environment.

Containers are also central to DevOps adoption. Using a container architecture, Red Hat OpenShift Container Platform lets you easily and quickly build, develop, and deploy applications on nearly any infrastructure, including OpenStack. Native integration with Kubernetes container orchestration automates deployment and scaling of containerized applications.

Running and managing containers at scale requires smart infrastructure, which Red Hat OpenStack Platform provides well. In fact, the combination of Red Hat OpenShift Container Platform and Red Hat OpenStack Platform is ideal for organizations that want to deploy, run, and manage containers in an on-premise private cloud environment. The resulting infrastructure is as flexible, automated, and scalable as the workloads and containers running on it. IT operations teams can provide developers with the services and tools needed for application development faster—from behind your firewall and within IT's control. Plus, increasing integration between the two platforms means that a growing number of operations will become simple, one-time actions. All of these capabilities help you rapidly bring new product and services to market, so your organization can stay competitive.



KUBERNETES

is the most popular tool for container management in OpenStack and is used in **50% of deployments.**

THE OPENSTACK FOUNDATION
USER SURVEY
NOVEMBER 2017

WHAT?

Run containerized applications on Red Hat OpenStack Platform and manage them efficiently through Kubernetes orchestration.

WHY?

Containers are central to DevOps and enable fast application development, deployment, and management. Portability lets you move workloads across environments as needs change.

HOW?

Deploy Red Hat OpenShift Container Platform on top of Red Hat OpenStack Platform to get started.

LEARN MORE

Visit these links to learn more:


- » [Red Hat OpenShift Container Platform documentation](#)
- » [Deploying and managing Red Hat OpenStack Container Platform 3.4 on Red Hat OpenStack Platform 10 reference architecture](#)
- » [Deploying and managing Red Hat OpenStack Container Platform 3.6 on Red Hat OpenStack Platform 10 reference architecture](#)

⁹ "OpenStack User Survey." The OpenStack Foundation. April 2017.
openstack.org/assets/survey/April2017SurveyReport.pdf

GENERAL RESOURCES

Where to find more information and guidance

Red Hat provides many resources for Red Hat OpenStack Platform. The majority of these resources are located on the [product documentation for Red Hat OpenStack Platform](#) page on the [customer portal](#). From this page, you can navigate to detailed documentation, knowledgebase articles, videos, and discussions. For a general overview of the latest version of Red Hat OpenStack Platform, we recommend reading the release notes and product guide, both located on the main documentation page.



Building a cloud infrastructure with OpenStack may seem like a daunting task, but Red Hat Consulting experts can help. All Red Hat Consulting engagements begin with a complimentary on-site discovery session. During these half-day sessions, Red Hat experts work with you to identify your business challenges and viable approaches for overcoming them, as well as necessary participants within your organization and desired outcomes for implementing emerging technologies.

**SCHEDULE A
COMPLIMENTARY
DISCOVERY SESSION:**
redhat.com/consulting




Copyright © 2018 Red Hat, Inc.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and Ansible are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. 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. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community.

F11834_0418_KVM



RED HAT
OPENSTACK
PLATFORM

