PLATFORM9

Deployment Models for OpenStack

Table of Contents

1. Executive Summary

OpenStack's phenomenal growth has made it the industry's leading private cloud management platform and the second largest open source project of all time. As seen in the graph below, search interest in OpenStack has far exceeded interest in alternative private cloud management platforms, including VMware's vCloud/ vRealize solution.

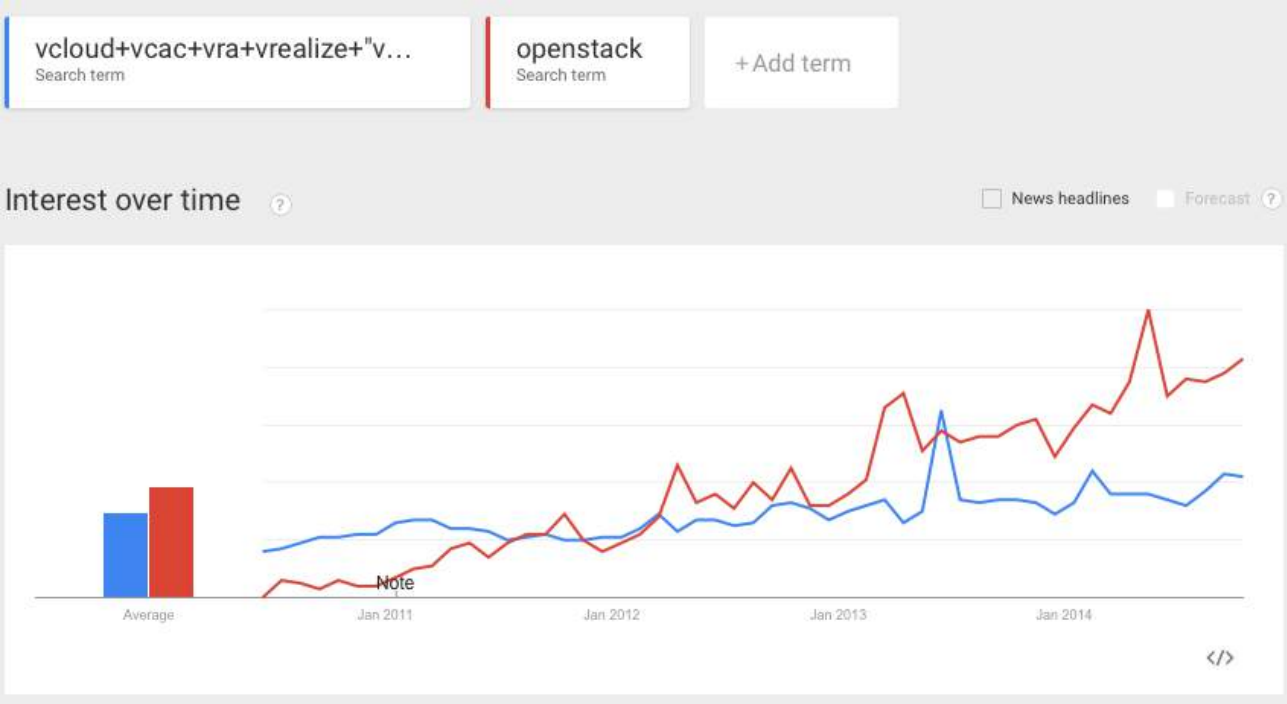

Figure 1: Search Interest reflects OpenStack momentum

Several OpenStack deployment models have emerged, each with its own pros and cons. Platform9 Managed OpenStack provides the simplest model for enterprises to implement a private cloud. Platform9 enables the use of existing or new virtualized servers to power private clouds in minutes, while interoperating seamlessly with existing processes and management solutions.

2. OpenStack Architecture Overview

OpenStack's design is inspired by Amazon Web Services (AWS), with well documented REST APIs that enable a self-service, elastic Infrastructure-as-a Service (IaaS) cloud. In addition, OpenStack is fundamentally agnostic to the underlying infrastructure, integrating with various compute, virtualization, network and storage technologies.

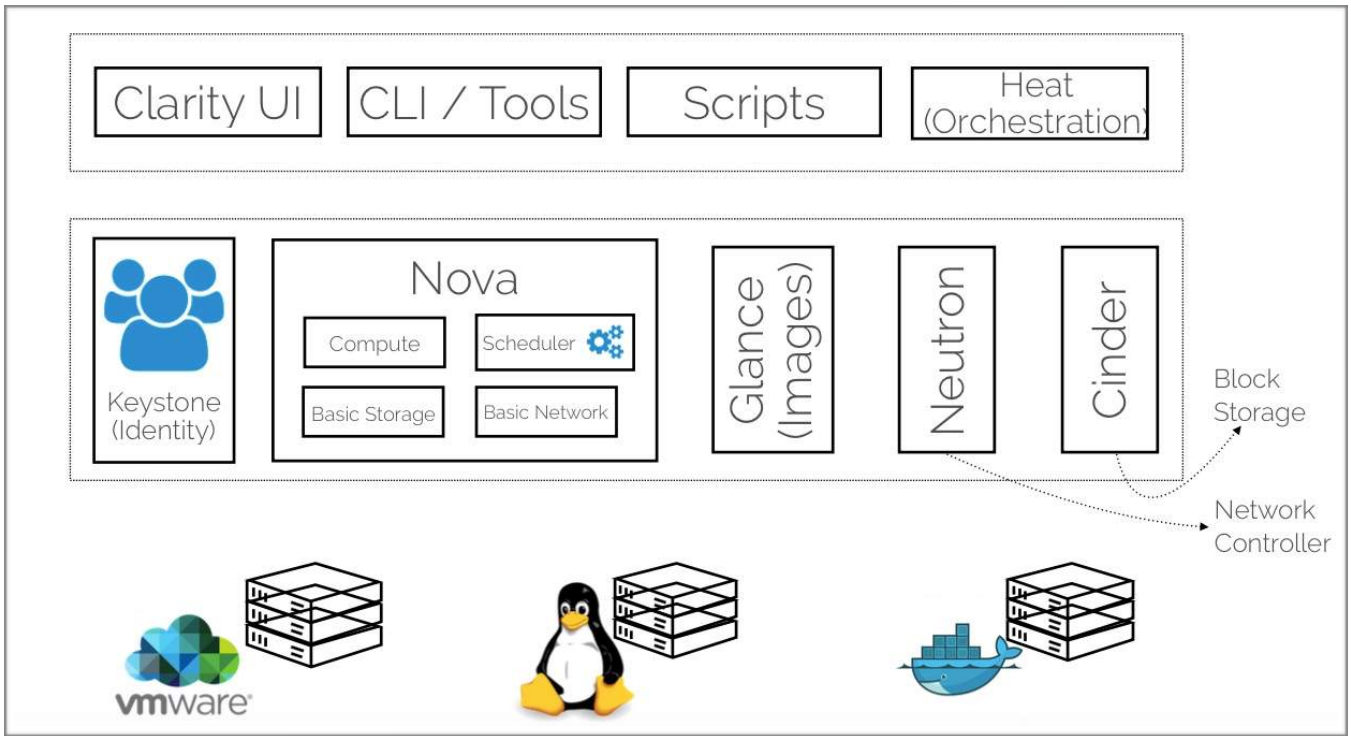

Figure 2: OpenStack architecture is loosely coupled, and extensible to support any Hypervisor / Container; Storage, and Network system

3. Mismatch of Expectations with OpenStack

OpenStack has multiple advantages over other private cloud platforms, such as simple REST APIs, an AWS-like service oriented architecture, and a management platform that works across multiple virtualization technologies. Despite these advantages, OpenStack adoption suffers from the following 3 limitations:

1. OpenStack TCO is Variable

Deploying OpenStack in production is a challenging and resource sensitive exercise. Analysts estimate that a team of at least 20 engineers would be required to deploy OpenStack for a 25 rack infrastructure[1]. Figure 3 shows that the actual cost of deploying OpenStack is much higher than the perceived cost of deploying the software.

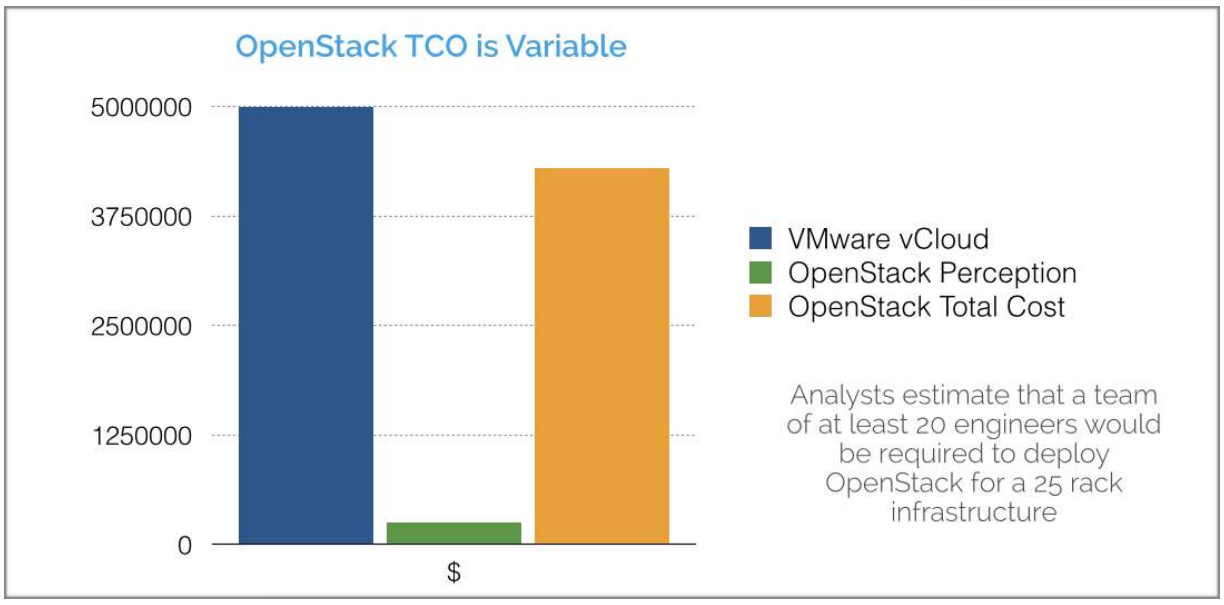

Figure 3: Perceived cost of deploying OpenStack << Actual cost of deploying OpenStack

[1] http://blogs.barrons.com/techtraderdaily/2014/04/03/vmware-price-tag-versus-openstack-immaterial-says-isi/?mod=yahoobarrons&ru=yahoo

2. Upgrade Issues

Upgrading from one version of OpenStack to another is a difficult task that can be plagued with unplanned downtimes arising from unexpected issues. Troubled by past upgrade processes many organizations are hesitant to migrate to the latest OpenStack release. This fact is supported by Figure 4, which displays the results of an OpenStack foundation survey[2]. According to the survey, 78% of organizations use an OpenStack release which is more than 2 years old.

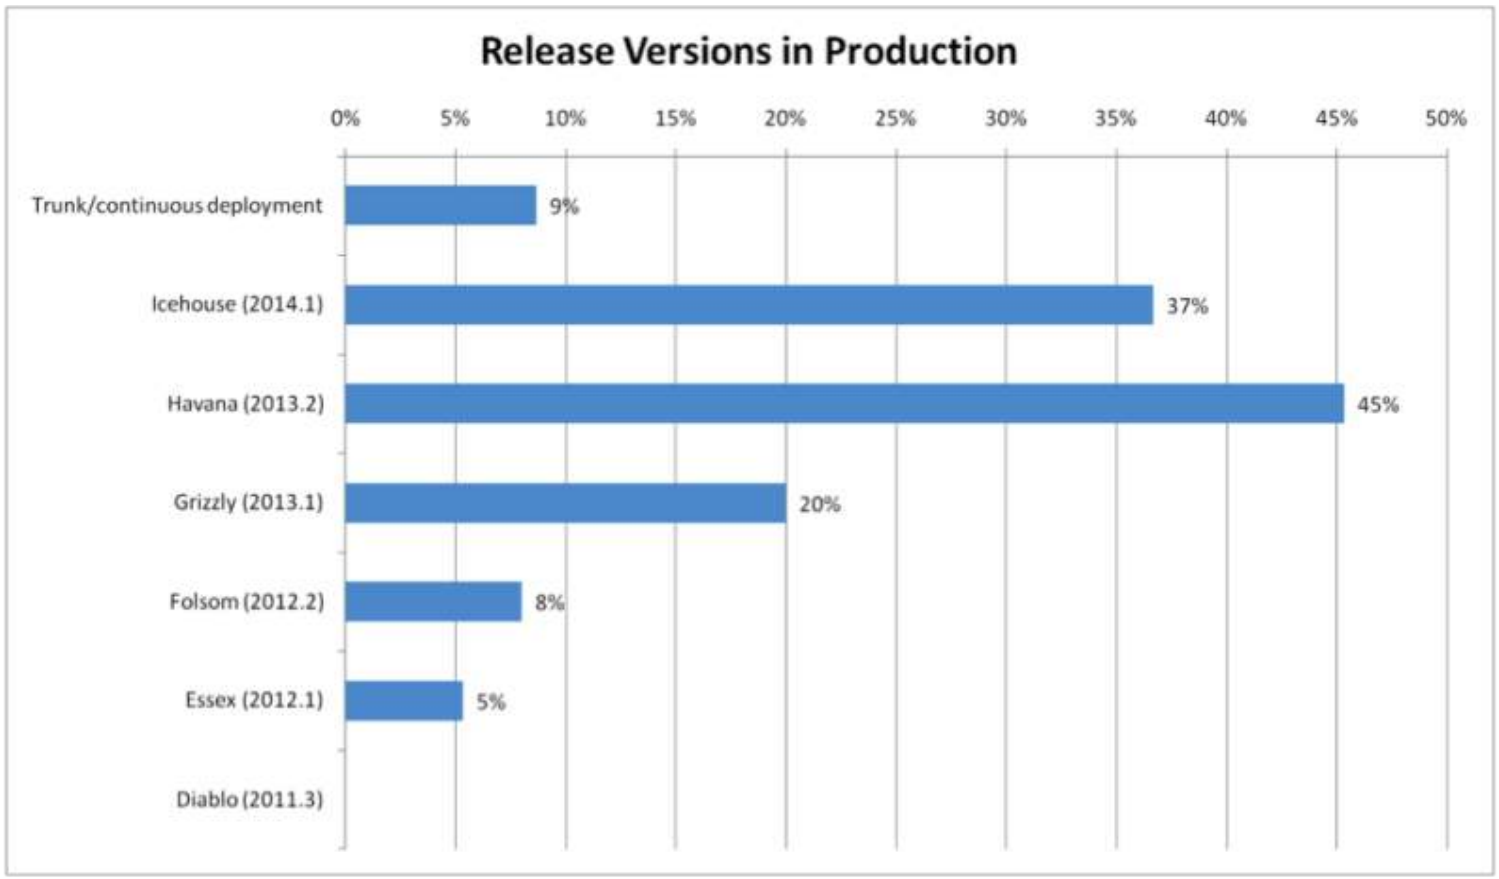

Figure 4: 78% of organizations use an older version of OpenStack to avoid upgrade issues

[2] http://superuser.openstack.org/articles/openstack-user-survey-insights-november-2014

3. Does Not Work Out-of-the-Box

OpenStack is an open source project, not a packaged and production ready product. This makes deploying OpenStack directly from Master an arduous and engineering-intensive task when compared with deploying a packaged product from a distribution vendor or consuming OpenStack as a service.

4. Deployment Models for OpenStack

1. On-Premises Distribution

This is the most commonly adopted deployment model. IT is responsible for installing, configuring, monitoring, and troubleshooting their cloud. This model uses in-house personnel to customize OpenStack to meet organizational requirements. This deployment model can be cost effective for organizations when implemented successfully.

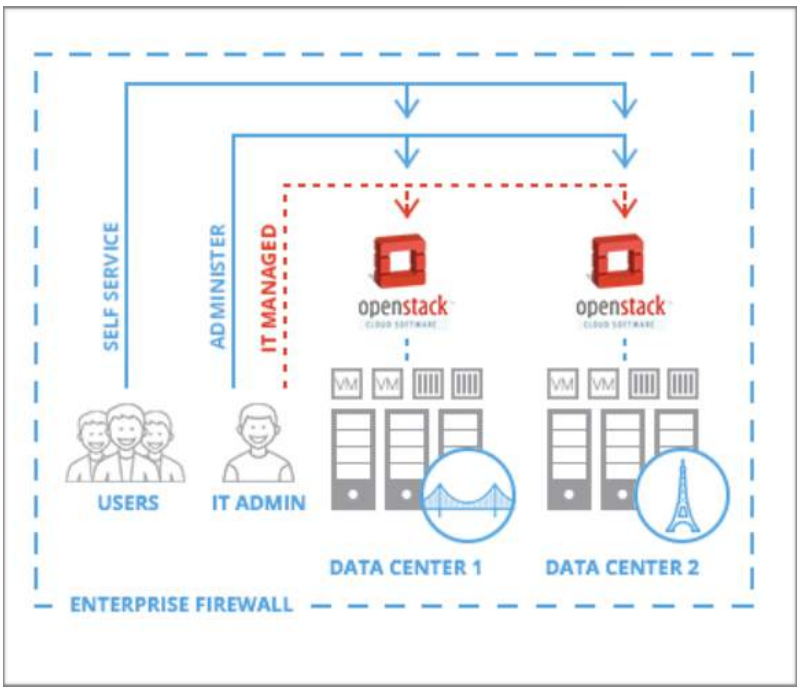

Figure 5: An on-premises OpenStack distribution uses in-house secure infrastructure

Pros

1. Secure: No workload data leaves your network perimeter
2. Customizable: OpenStack can be customized to suit an organization's requirements

Cons

1. High Operational Costs: There are high costs associated with configuring, monitoring, backing up, and upgrading OpenStack
2. Creates silos: Having different OpenStack deployments across multiple geographies can create infrastructure silos. This can create redundancy and complexity that is difficult to manage

2. Hosted Private Cloud

This model of deploying OpenStack involves using a service provider to deploy a hosted private cloud in the service provider's data centers. The service provider is responsible for installing, configuring, monitoring, upgrading and troubleshooting OpenStack as well as the customer's infrastructure. This contractual model is usually based on SLAs and makes it simpler for organizations because they able to leverage the service provider's data centers, hardware, and OpenStack expertise.

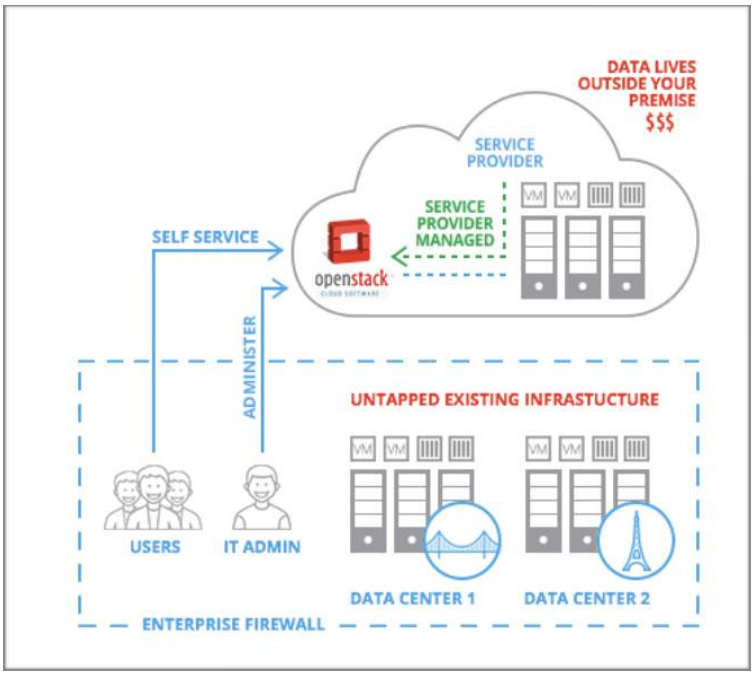

Figure 6: In a hosted private cloud model, data lives outside the organization's premises.

Hosted Private Clouds are simpler to deploy and as a result there is a higher adoption. This seems to be supported in the below survey by the OpenStack Foundation, conducted in November 2014.[3]

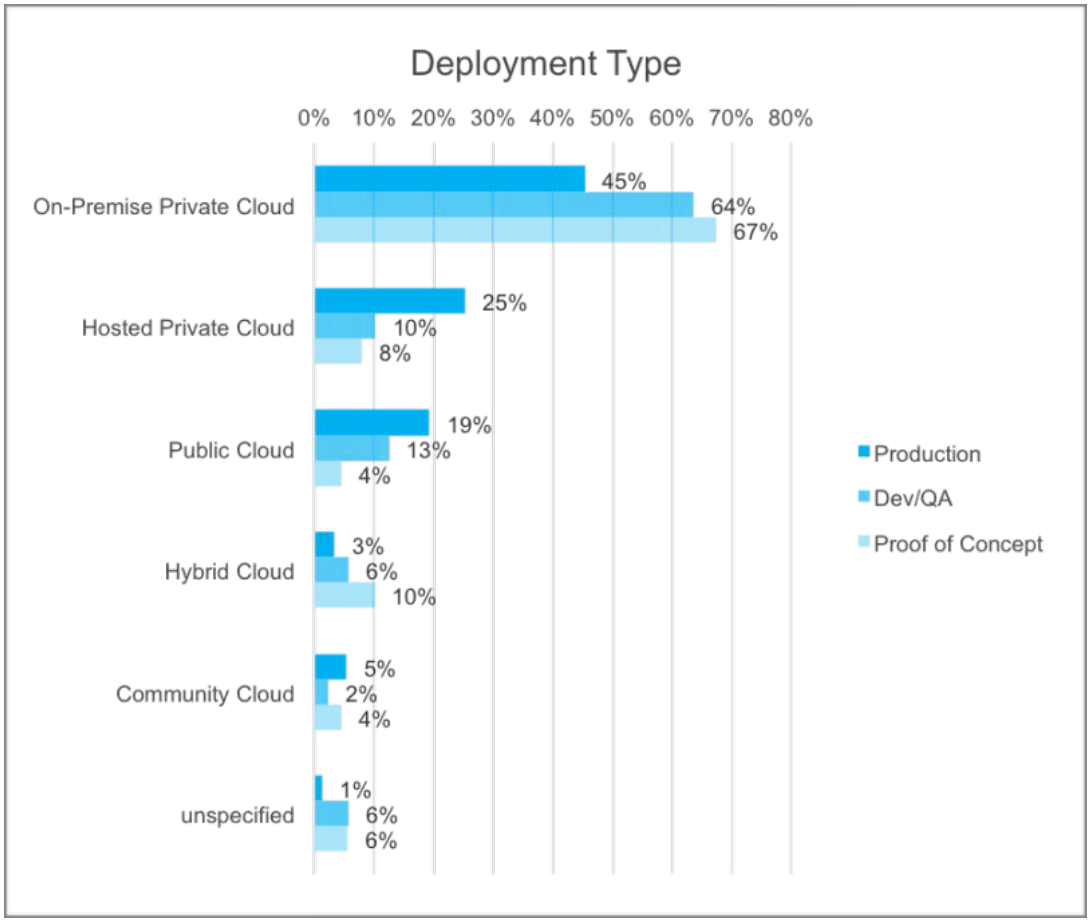

Figure 7: Hosted Private Clouds are favored for a production environment.

Pros

1. Convenience: Organizations who want a dedicated private cloud without owning the hardware can conveniently deploy one with ease
2. Simplicity: This model saves organizations from the operational complexity of OpenStack

Cons

[3] http://superuser.openstack.org/articles/openstack-user-survey-insights-november-2014

1. Creates silos: Existing infrastructure might be unused. This can lead to fragmentation and redundancy
2. Security: Workload data leaves your network perimeter. Organizations depend on the service provider to build a secure and operational OpenStack

3. Cloud Management-as-a-Service

Platform9's unique Cloud Management-as-a-Service approach combines the advantages of the on-premises OpenStack distribution and the hosted private cloud models. The Cloud Management-as-a-Service model provides the economics, data locality, and infrastructure choice of an on-premises deployment model with the convenience and operational efficiency of hosted private clouds. Organizations that use Platform9 Managed OpenStack host their data on their own infrastructure while Platform9 hosts and manages the OpenStack control plane.

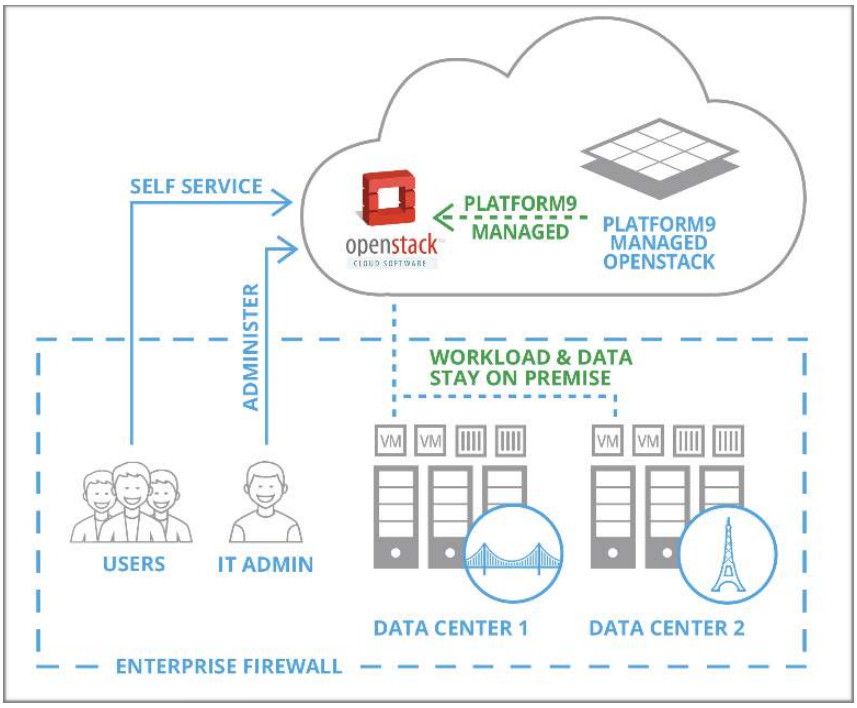

Figure 8: Platform9's Cloud Management-as-a-Service hosts the OpenStack control plane and leverages on-premises infrastructure

Pros

1. Set up in minutes: IT Team Members can build an OpenStack private cloud in minutes
2. Eliminates silos: Platform9 provides a single pane of glass across geographies and different virtualization platforms such as KVM, VMware vSphere, and Docker
3. Secure: No workload data leaves the user's network perimeter
4. Greenfield and brownfield: Platform9 works seamlessly with existing or new infrastructure

Cons

1. HTTPs access: This model requires outbound, secure HTTPS access from the organization's servers to the Platform9 OpenStack controller

5. Key Takeaways

This white paper discussed three different models for deploying OpenStack. Working with multiple customers, Platform9 has found that organizations adopting some type of OpenStack-as-a-Service model and particularly an approach such as Platform9 Managed OpenStack, realize significant cost and time savings when deploying a private cloud.

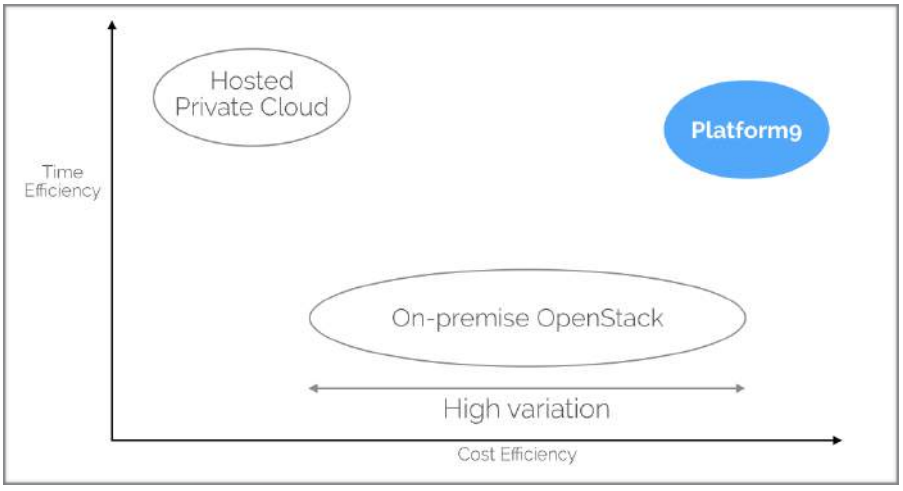

Figure 9: Platform9's OpenStack-as-a-Service approach saves money and time for organizations deploying a private cloud

About Platform9

Platform9 makes private clouds easy by delivering OpenStack-as-a-Service. To learn more about Platform9's OpenStack-as-a-Service, request your free self-service trial at www.platform9.com. You can also email us at support@platform9.com or call at +1-650-898-7369.