A Tale of Two Workloads

USING OPENSTACK® AND VMWARE VSPHERE® TO BUILD A HYBRID SOLUTION

MANAGED VIRTUALIZATION
Powered By VMware®

CONTENTS

A Tale of Two Workloads

**USING OPENSTACK® AND VMWARE VSPHERE®
TO BUILD A HYBRID SOLUTION**

EXECUTIVE SUMMARY

With the recent release of Icehouse, the latest version of the OpenStack cloud software, OpenStack is continuing to mature as an enterprise-grade platform that is being used to power both public and private clouds. Rackspace, as the founder of OpenStack, uses it to power the largest OpenStack-powered public cloud and a leading OpenStack-powered private cloud offering.[1]

As cloud computing continues to grow in adoption, many companies are evaluating OpenStack and trying to determine how it should fit into their overall IT strategies for private cloud, public cloud and hybrid clouds. Many of these companies are current VMware customers who are also trying to understand where their vSphere-based virtualized infrastructure fits into their overall cloud and data center strategy.

Key to this transition to a cloud computing model, for customers, is understanding the requirements of their application workloads and being able to map them to the correct infrastructure so that they can gain maximum agility and efficiency without sacrificing application availability. Otherwise, they may find themselves using the wrong "cloud" as the foundation for their present and future mission-critical workloads.

UNDERSTANDING THE CHALLENGES

Rackspace has helped over 100 customers in the implementation of their OpenStack-powered private clouds. Along the way, we have spoken with a number of companies that are currently conducting evaluations of OpenStack or using it in small projects. Almost all of these companies are current VMware customers with legacy application workloads. They understand that OpenStack is not a technology they can afford to ignore, but they often struggle to understand its true value and how it should impact their current vSphere deployments. The questions they frequently ask include:

• Is OpenStack a free alternative that can be used to replace our current VMware vSphere infrastructure?

• Is there feature parity between OpenStack and VMware vSphere, i.e. High-Availability, vMotion, Distributed Resource Scheduler, etc.?

• Can we run the same application workloads effectively on both platforms?

• Are there application workloads that more suitable for one platform over the other?

• Can we and should we use OpenStack together with VMware vSphere?

Move forward in the new world of cloud-native scale-out applications while continuing to effectively manage legacy applications.

Answering these questions is crucial, for current VMware customers, to be able develop a strategy that allows them to move their organization forward in the new world of cloud-native scale-out applications while continuing to effectively manage their legacy applications.

DEVELOPMENT OF A TWO-WORKLOAD STRATEGY

The starting point for customers is to gain an understanding of these two workload types and then to classify all the applications in their current and planned portfolio. This will allow them to determine if they should adopt an OpenStack cloud and/or if they should continue running a VMware vSphere infrastructure. It will also allow them, if they choose to run both environments, to determine what workloads should be hosted by each infrastructure.

Rackspace recommends that customers begin determining their two-workload strategy by answering the following questions:

• What are the differences between legacy and cloud-native workloads and how are they managed differently?

• What are the various legacy infrastructure and cloud consumption models in the market today?

• What are the viable options for using VMware vSphere together with OpenStack?

UNDERSTANDING THE TWO WORKLOADS

Application requirements should dictate infrastructure. This fundamental principle means that we need to know what our present and future application requirements are to determine if they should sit on a virtualized infrastructure, such as VMware vSphere, or a cloud platform, such as OpenStack. First, let's focus on legacy applications that currently run on virtualized infrastructures, such as the VMware® ESXi™ hypervisor and VMware vSphere. ESX/ESXi and vSphere came into prominence as a technology to virtualize many smaller servers so they can be consolidated on to a few large servers. This worked very well since most servers were hosting applications with monolithic architectures, such as Oracle or Microsoft Exchange. Today, each instance of this type of legacy application is still typically encapsulated in a single virtual machine and grows by scaling up on a single physical server running the ESXi hypervisor. High availability can be achieved by running a clustered version of the application, such as Oracle Real Application Clusters; however, this can be an expensive and overly complex solution and most applications do not have such functionality. Most VMware shops choose to run their application servers as virtual machines in vSphere clusters and depend on features such as vSphere HA and vMotion to provide infrastructure resiliency and redundancy. While these solutions work well, they also require certain architecture choices

> By moving application resiliency up the stack, cloud platforms remove the need for shared storage.

to be made, such as reliance on shared storage, which makes on-demand rapid scaling-out of theses infrastructures difficult.

Cloud computing, however, was created to accommodate a different class of cloud-native applications, such as mobile and web software. A few cloud platforms like OpenStack are designed to host applications that have a distributed architecture where application components are distributed across multiple physical or virtual servers. These applications are generally designed to grow by scaling out across multiple servers so as demand increases, resources can be expanded by adding more application instances and re-balancing workloads across those instances. Another design principle behind cloud platforms such as OpenStack is that given the distributed nature of these applications, ownership for resiliency belongs to the applications and not the underlying infrastructure. This approach is often misunderstood by folks in the VMware space as a shortcoming and immaturity in the OpenStack platform; "lacking" features such as vSphere HA are seen as a potential warning sign that OpenStack is not an enterprise-grade platform.

However, this is a misunderstanding of the different design principles behind legacy and cloud applications and infrastructures. Distributed applications that run on cloud platforms have lowered the barrier for building resiliency, both in terms of cost and of usability. By moving application resiliency up the stack, cloud platforms remove the need for shared-everything architecture based decisions such as the use of shared storage. This promotes the use of commodity as an option for running a cloud platform and creates an architecture that enables rapid scaling out the infrastructure. It is also an architecture that is best suited for next-generation large-scale application environments where failure is expected and needs to be designed at multiple layers, not simply at the infrastructure layer.

UNDERSTANDING THE CONSUMPTION MODELS

Once a two-workload strategy has been created with a careful mapping of every workload to the appropriate infrastructure, the next step is to understand the different infrastructure consumption models. Here, it is important to differentiate between virtualized infrastructure consumption and cloud consumption in order to choose the appropriate consumption model for your workloads.

Along with operating their own on-premises virtualized infrastructures, such as vSphere, companies can consume managed hosting offerings such as the Rackspace Dedicated VMware® vCenter Server™. This is built on VMware technologies and offers off-premises virtualization solutions to augment customers' on-premises vSphere deployments. It is designed and ideally suited for legacy applications that do not require rapid scaling and rely on the virtualized infrastructure to provide application availability and resiliency. This off-premises offering provides value by allowing customers to expand

3 CONSUMPTION MODEL OPTIONS TO RUN BOTH WORKLOADS:

1 Infrastructure Silos

2 Multi-Hypervisor Integration

3 Best-fit Hybrid Solution

their legacy workload footprint and to setup a disaster discovery solution without incurring high CapEx costs. It may also provide the added benefit of minimizing the change in operational processes since customers can leverage the same tools and methodology for both their on-premises and off-premises infrastructures.

In contrast, cloud consumption typically begins with public cloud usage such as AWS or the Rackspace Cloud and may later include private cloud deployments such as the Rackspace Private Cloud. In this space, the focus is on accommodating next-generation applications and being able to scale and to provision resources quickly, often using commodity hardware. These cloud platforms are not typically suited for legacy workloads, unless some integration work is done with an infrastructure that is designed to be resilient. There is a steeper learning and adoption curve as operators will need to learn new tools and processes for managing their cloud environment.
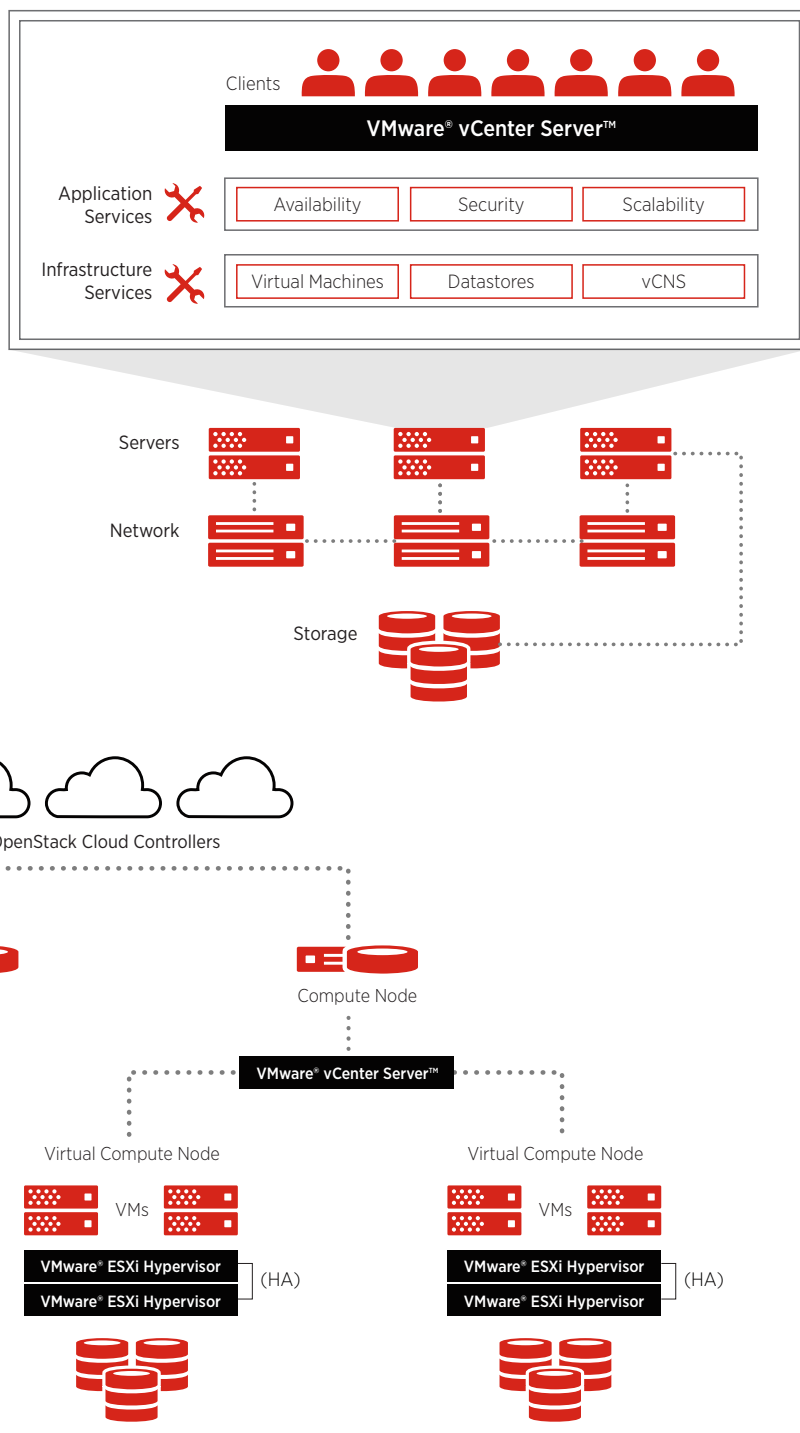
UNDERSTANDING VSPHERE WITH OPENSTACK

It should be clear by now that one size does NOT fit all when it comes to building out infrastructure for different workloads. Rackspace has customers that, because they are already running their business on cloud-native type applications, have been able to move directly to our OpenStack-powered public cloud and/or our private clouds. However, most customers have legacy applications, often running on virtualized infrastructures that cannot be easily rewritten to use a cloud platform such as OpenStack. For these customers, co-existence and not replacement is the route they will need to take in adopting OpenStack into their portfolio. For these customers, the next step to take after classifying their applications and choosing their consumption models is to decide how to run both workloads in their environment. This route typically fork into one of three paths:

INFRASTRUCTURE SILOS

This is the route most frequently chosen by customers, though not necessarily the best route. Typically, this involves making the decision to keep legacy applications running on a vSphere infrastructure while building new applications on a separate OpenStack cloud. While this is the least disruptive route to take in adopting OpenStack, it also perpetuates IT silos and adds complexity and the additional overhead of managing two completely distinct environments since different operational tools and processes are needed and separate operations teams are often required.

VMware vSphere

Clients

VMware® vCenter Server™

Application Services

| Availability | Security | Scalability |

Infrastructure Services

| Virtual Machines | Datastores | vCNS |

Servers

Network

Storage

OpenStack Cloud Controllers

Compute Node / KVM

Compute Nodes / KVM

Compute Node

Storage

Storage

VMware® vCenter Server™

Virtual Compute Node

Virtual Compute Node

VMs

VMs

VMware® ESXi Hypervisor

VMware® ESXi Hypervisor

(HA)

VMware® ESXi Hypervisor

VMware® ESXi Hypervisor

(HA)

MULTI-HYPERVISOR INTEGRATION

Another possible route is to leverage the work VMware has done to integrate vSphere into OpenStack. This is similar to the silos route where legacy workloads continue running on vSphere while next-generation workloads can run on a hypervisor such as KVM or Xen. However, OpenStack is used here

as the control plane to manage a multi-hypervisor cloud, consolidating cloud management while allowing applications to be hosted on the environment best suited for them. Since the OpenStack management plane communicates with vCenter for vSphere management, applications running on vSphere are able to continue leveraging infrastructure resiliency features such as HA and vMotion. Cloud-native applications, as mentioned earlier, run in their own non-vSphere "hypervisor zone" but managed by the same top-level control plane.

The primary drawback to this route is that the vSphere integration is relatively new and not all functionality are tightly integrated with OpenStack. Also, because of the architectural decision, made by VMware, to have OpenStack manage vSphere using vCenter as the proxy management node, provisioning can be performed via OpenStack but configuration of the vSphere clusters must be performed separately using vCenter and the vSphere Web Client.

BEST-FIT HYBRID SOLUTION

A third route offered by Rackspace is to leverage a best-fit hybrid solution. This looks initially like the silos approach in terms of creating separate infrastructures for vSphere and OpenStack to help ensure that workloads are hosted on the infrastructure that best fits their needs. However, the goal here is not merely to maintain separate infrastructures, but to integrate them together into a true hybrid solution.
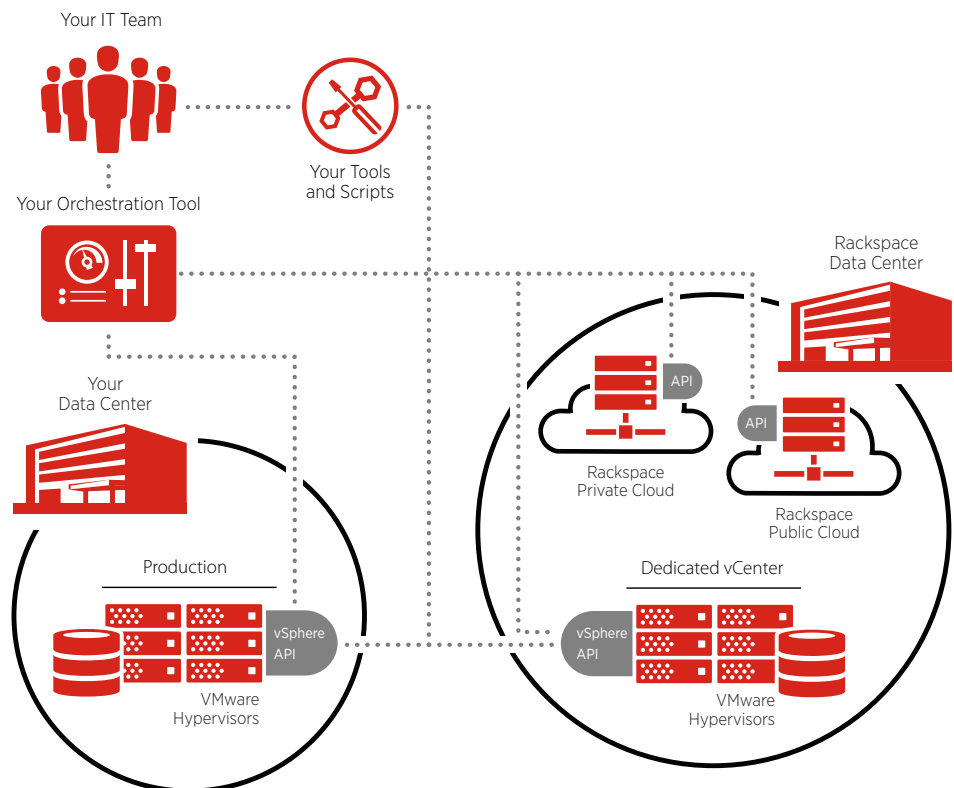
This is accomplished by integrating the data plane and the operational processes of these separate infrastructures. By connecting them at the network layer, we enable the sharing of data and integration of applications across all infrastructures, regardless of consumption model. A representative use case is a multi-tier application with a distributed web-tier running on an OpenStack-powered Rackspace Private Cloud that is connected, via Rackspace's RackConnect solution, to an Oracle database running on a managed vSphere cluster. In this hybrid solution, the Rackspace Public Cloud can also be used to enable bursting up of web-tier workloads from the private cloud during peak seasons.

CONCLUSION

The push to cloud computing is only accelerating and the Business is demanding new levels of responsiveness and agility from their IT departments. The question isn't "should we adopt cloud," but "when to adopt cloud" and "what is the best approach to take." In addition, IT must take into account their existing environments and continue to support them adequately while building new infrastructures to support the next generation of mobile and web applications.

The way forward for enterprises is to evaluate and to classify their current and planned application workloads. Enterprises then need to map these workloads to the correct infrastructures and if multiple infrastructures are required, decide how to best manage and integrate these separate infrastructures.

Rackspace has created the consumption models and the expertise required to help customers to develop and to enable their two-workload strategy. As one of VMware's largest VSPP partners, Rackspace has expert VMware Certified Professionals available and experience that comes with managing over 45,000 VMs. To find out more about how Rackspace can help you in your adoption of cloud computing, please call 800-961-2888.

ABOUT RACKSPACE

Rackspace® (NYSE: RAX) is the global leader in hybrid cloud and founder of OpenStack®, the open-source operating system for the cloud. Hundreds of thousands of customers look to Rackspace to deliver the best-fit infrastructure for their IT needs, leveraging a product portfolio that allows workloads to run where they perform best—whether on the public cloud, private cloud, dedicated servers, or a combination of platforms. The company's award-winning Fanatical Support® helps customers successfully architect, deploy and run their most critical applications. Headquartered in San Antonio, TX, Rackspace operates data centers on four continents. Rackspace is featured on Fortune's list of 100 Best Companies to Work For.

For more information, visit www.rackspace.com

NOTES

1 : Rackspace UK Inc., http://www.rackspace.co.uk/open.

Toll Free: 1.800.961.2888 | International: 1.210.312.4700 | Or Visit Us www.rackspace.com

Rackspace® | 1 Fanatical Place | City of Windcrest, San Antonio, Texas 78218 U.S.A | DATE MODIFIED: 10272014