openstack
CLOUD SOFTWARE

Adding Speed and Agility to Virtualized Infrastructure with OpenStack

OPENSTACK WHITE PAPER

Contributors:

Amr Abdelrazik, *Manager, Product Marketing*, OpenStack Cloud Platform, VMware
Greg Bunce, *Intel IT Infrastructure Engineering Team*, Intel Corporation
Kathy Cacciatore, *Consulting Marketing Manager*, OpenStack Foundation
Kenneth Hui, *Business Development Manager*, OpenStack-Powered Cloud Solutions, EMC Corporation
Sridhar Mahankali, *Intel IT Infrastructure Engineering Team*, Intel Corporation
Frans Van Rooyen, *Compute Platform Architect*, Adobe Digital Marketing

Executive Summary

Today's software-defined economy requires businesses to move faster than their competitors. Speed and agility is critical to keeping up with competitive demands for new applications and big data analytics, as well maintaining existing infrastructure.

Many enterprise organizations and their software developers are turning to OpenStack® for its open APIs, flexible architecture and large commercial ecosystem to compete in a completely new paradigm of software development and deployment.

> According to Gartner, "The modular architecture of OpenStack provides extreme flexibility in terms of implementation, particularly when compared with the fixed-function design of many monolithic commercial CMPs."[1]

Many of these enterprises also operate large VMware-based virtualized infrastructures to support a legacy of mission-critical, scale-up applications and enterprise databases. They value VMware for its compute, network, storage and management technologies. But they are also encountering new use cases that demand agility.

Can they combine the agility offered by cloud computing while maintaining the stability and resiliency of their traditional environment? The simple answer is, yes!

When considering OpenStack, VMware vSphere® customers often ask three questions. Is OpenStack a free and open-sourced hypervisor that can be used to replace our current VMware ESX® servers? Is there feature parity between OpenStack and vSphere? Are there benefits to using both?

This white paper addresses these questions by understanding the differences between legacy and cloud workloads; and detailing three approaches to using VMware technology with OpenStack.

Virtualization and Cloud

To understand the differences between traditional infrastructures and cloud computing, look at how they differ in the philosophy of their designs.

First, consider the infrastructures that were built on virtualization technologies, such as the VMware ESXi™ hypervisor and vSphere. These technologies offered consolidation on a smaller number of larger servers.

The solutions worked well because most servers hosted applications with monolithic architectures, such as Oracle or Microsoft® Exchange. Today, each instance of this type of application is still encapsulated in a single virtual machine and grows by scaling up on a single physical server running the ESXi hypervisor.

To provide these legacy applications with a resilient infrastructure, most VMware shops choose to run their application servers as virtual machines in vSphere clusters, depending on proven features such as VMware vSphere® High Availability and VMware vSphere® vMotion®. While these solutions work well, they also require certain architectural choices be made, such as reliance on shared storage that makes scaling out difficult.

Next, consider the difference in cloud computing. Cloud platforms like OpenStack are designed to be used with a different class of applications, such as MongoDB and Hadoop that are architected to scale horizontally, and are resilient against virtual machines shutdowns. Resources can be expanded by adding more application instances and re-balancing workloads across those instances. These distributed applications are responsible for their own resiliency, independent of the underlying infrastructure and advanced hypervisor features.

VMware users often mistake this as a reflection of deficiencies and immaturity of the OpenStack platform. However, this assertion is based on a misunderstanding about the differing design principles of legacy applications and cloud computing.

By moving application resiliency up the stack, cloud platforms remove the need shared-everything architecture-based decisions. Commodity hardware is seen as an option for running a cloud platform and creates an architecture that enables rapid scaling of the infrastructure.

This architecture is best suited for next-generation, large-scale application environments where failure is anticipated and requires a design at multiple layers, beyond the infrastructure.

OpenStack Overview

OpenStack provides the tools and technologies to abstract the underlying infrastructure in an easy and standardized consumption model. OpenStack sits on virtual or physical Compute, Network and Storage technologies and provides the APIs and tools to access these resources in an agile and programmatic manner.

OpenStack interfaces with the underlying infrastructure through open source or vendor-provided drivers. This standardized abstraction helps prevent customers from being locked into a specific technology or tool.

OpenStack provides additional services such as identity management, orchestration, and metering accessed in the same programmatic manner through the API. OpenStack also offers a framework for evolution to DevOps, Continuous Integration and Continuous Deployment methodologies.

OpenStack is not a hypervisor, but it supports several hypervisors through an abstraction layer. The software supports popular commercial and open source hypervisors including vSphere ESXi as well as KVM, Xen, QEMU and Hyper-V.

Cloud-native and traditional application developers can leverage the flexibility and agility of cloud application provisioning through APIs. OpenStack can streamline the application lifecycle from development, to unit and system testing, and production deployment while exploiting VMware vSphere features such as HA and vMotion. Traditional application teams can experiment with next generation application architectures and gradually transition their application and operation to a cloud native architecture.

High availability can be architected into an OpenStack infrastructure. In next-generation infrastructures, failure handling is designed into the applications to reduce hardware implementation costs. In contrast, an infrastructure-level, high-availability architecture may require redundant hardware, storage and networking, third party software and custom code. The OpenStack development community has documented how to architect passive/active or active/active scenarios in the **OpenStack High Availability Guide.**

Live migration of virtual machines is under development by the OpenStack community.

Using vSphere with OpenStack

Companies with distributed applications can move directly to OpenStack-powered public or private clouds. However, most customers have legacy applications. These applications often run on a bare-metal or virtualized infrastructure. They are not easily rewritten for a cloud platform such as OpenStack.

In addition, enterprises with VMware virtualized servers typically invest largely in license, skillsets and integrations with management tools. With OpenStack on top of VMware, these resources can be leveraged while providing the capability to build cloud native applications, and gradually transition the existing application development model to the cloud.

For these customers, co-existence of the two platforms is the default way to adopt OpenStack within their current vSphere data centers. Whereas virtualization adoption was driven by central IT, cloud computing is generally driven by developers. Two separate infrastructures can be maintained - vSphere and OpenStack. However, there are better ways to maintain both platforms.

vSphere support and deeper integration approaches are enabled by VMware contributions that integrate vSphere, VMware NSXTM and VMware Virtual SAN™ with OpenStack. Many VMware partners are OpenStack members who offer integrated products that make possible several OpenStack-VMware architectures.

In the next section we explore the three most typical architectures for deploying OpenStack with VMware.

Architectures for Deploying OpenStack with VMware

Using OpenStack to Manage a Multi-Hypervisor Environment

In this scenario, OpenStack is used as the control plane to manage a multi-hypervisor cloud, running both vSphere and alternate hypervisors such as KVM or Xen. This approach provides common, self-service provisioning and API access. It consolidates cloud management while allowing applications to be hosted on the environment best suited for them. For example, users can run a three-tier legacy application using a database backend such as Oracle on vSphere, and web-scale applications, such as MongoDB on KVM.
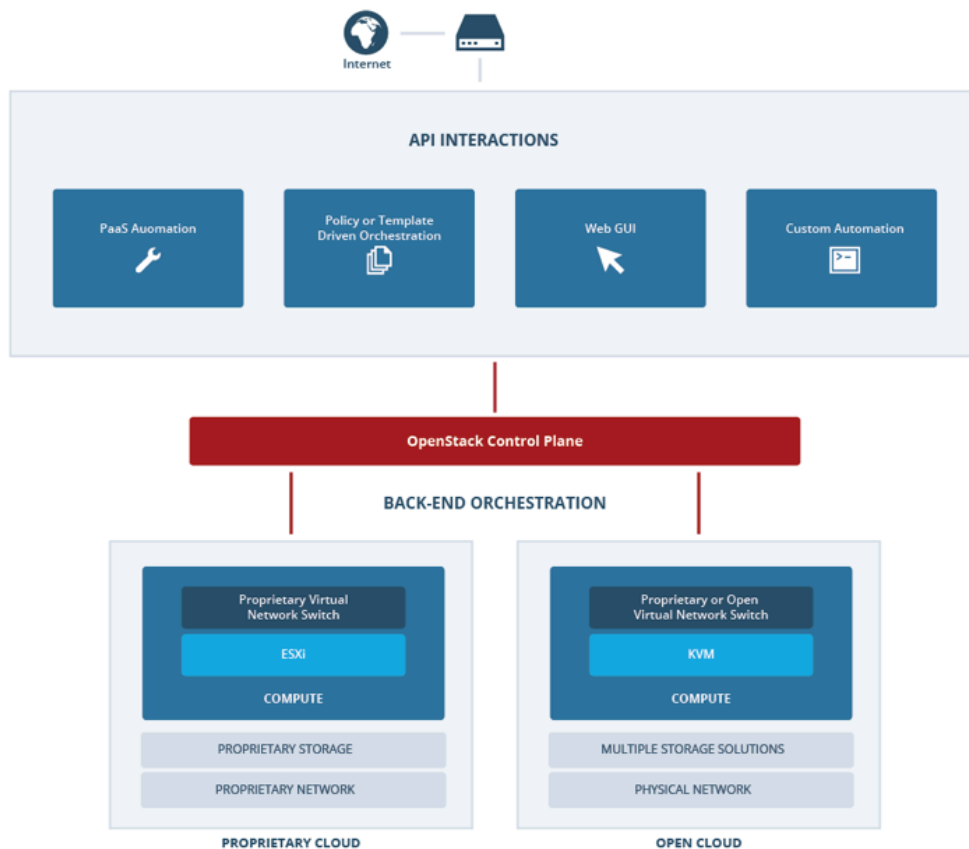

Figure 1: Example Multi-Hypervisor Cloud Architecture

The example architecture in Figure 1 supports multiple hypervisors, multiple storage solutions and virtual network integration. It shifts from IT managed to self-service and reduces the impact of migrating applications to new environments. A consistent user experience is provided while offering IT the flexibility to efficiently manage the overall infrastructure and resource utilization.

Intel® IT's hosting organization runs a large enterprise private cloud supporting mostly traditional enterprise workloads such as ERP, a mix of custom in-house developed applications, and commercial off-the-shelf (COTS) applications. In 2010, they implemented a custom private cloud - a VMware-based virtualization of the data center server environment with self-service capabilities - built from existing available components. Hosting operated a separate OpenStack cloud servicing a greenfield implementation of KVM and Ceph for provisioning both internal and externally-facing workloads.

The second and current instantiation of their private cloud is built upon OpenStack as the datacenter control plane to provide an abstraction layer to the legacy cloud infrastructure. This allows Intel to end-of-life the custom automation built for the initial private cloud.
By the end of 2015, they plan to be fully migrated to the common OpenStack control plane architecture depicted in Figure 1. Intel IT hosting is currently running this architecture at scale with a healthy mix of development and production workloads, exercising 17,000+ full create / use / destroy lifecycles over the last year: 7,000+ KVM VMs and 10,000+ ESXi VMs.

In the 2014 calendar year, Hosting fielded over 8,000 manual service requests accounting for approximately 190,000 hours spent awaiting fulfillment. Through proliferation of our OpenStack control plane architecture in 2015, they foresee elimination of 85% of manual service requests through instant service requests fulfilled through automation.

By the end of 2016 the objective is for 90% of routine service requests fulfilled instantly. The enterprise private cloud will be based primarily based on open standards and open source technologies. It represents the next step in the journey to a federated, interoperable, and open hybrid cloud.

A variation of this architecture is a multi-distribution/multi-hypervisor. It is applicable for organizations that want to start with one distribution/hypervisor and add other combinations in the future.

This architecture leverages the OpenStack concept of regions. Each hypervisor's complete OpenStack controller is in a region. Both regions are accessible to the developers through the standard API, and both regions leverage the same identity service (Keystone). Developers are afforded a seamless experience while IT has the flexibility to add the hypervisor of their choice.
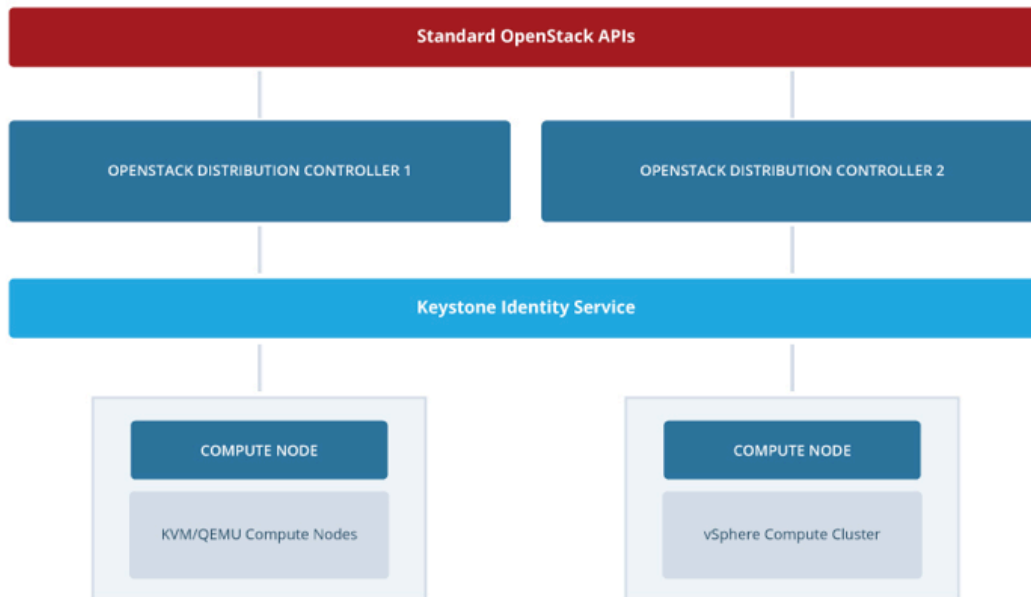
Figure 2: Example of Multi-Distribution/Multi-Hypervisor Architecture

OpenStack on vSphere (single hypervisor) using Commercial Distributions

Enterprises can also use VMware technologies with OpenStack by implementing a commercial OpenStack distribution.

The following OpenStack distributions support the VMware ESXi hypervisor in a cloud environment. They offer OpenStack CLI and REST API control of VMware compute, storage and networking resources for simplified, more agile cloud automation while retaining use of VMware tools like vMotion. Each OpenStack ecosystem member and VMware partner offers value-added services and support that complement and enhance OpenStack to meet differing enterprise needs.

- HP Helion OpenStack
- IBM Cloud Manager with OpenStack
- Mirantis OpenStack Private Cloud Software
- Platform9 Managed OpenStack
- Red Hat Enterprise Linux® OpenStack Platform
- SUSE OpenStack Cloud
- Ubuntu OpenStack
- VMware Integrated OpenStack

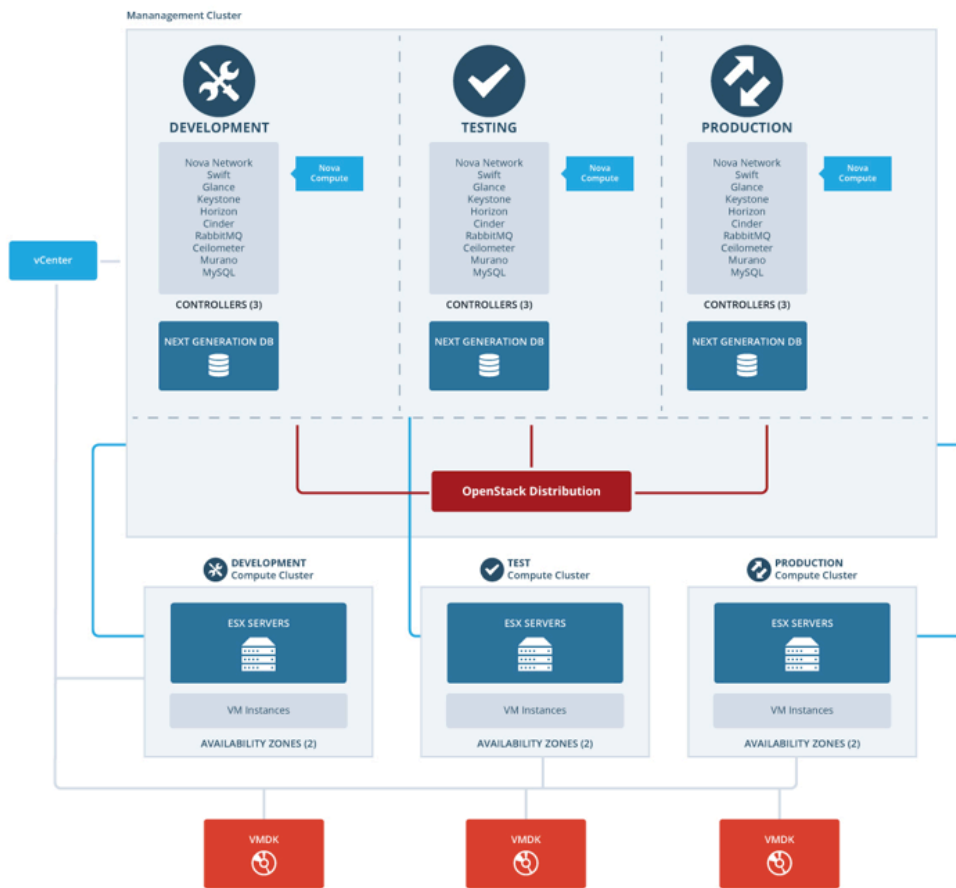Piston CloudOS+Piston OpenStack supports VMware NSX for software defined networking (SDN).

Figure 3: Example OpenStack Distribution Supporting VMware Infrastructure

08

The top section of Figure 2 is the management cluster or control plane. The boxes at the bottom are Glance images and the instances themselves. Images are ported into Glance as a VMDK file, which sits on the datastore. VMDK is an open file format that describes containers for virtual hard disk drives to be used in virtual machines like VMware ESXi.

IT professionals at a large multinational banking and financial company virtualized 70-80 percent of its servers with VMware and built a tremendous amount of expertise. They wanted the self-service capabilities of the cloud to provision ESXi virtual machines using VMware HA and vMotion without re-working existing monitoring and disaster recovery.

They decided on an OpenStack management cluster to leverage the existing infrastructure and expertise. They selected one of the commercial distributions supporting ESXi. The distribution vendor also provided consulting and implementation services, and ongoing support.

Today, the company's private cloud is used by application development, test and operations teams and has the capacity for thousands of internal applications. They run OpenStack in four datacenters with three clouds each (Dev, Test, Prod per Figure 2). They roll it out to additional datacenters regularly. Each user type – developer, tester and deployment operators - access only the cloud they need. vSphere HA and vMotion provide an extra level of built-in resiliency for these mission-critical applications.

OpenStack on vSphere (single hypervisor) using VMware Integrated OpenStack

The VMware Integrated OpenStack (VIO) distribution combines OpenStack and vSphere (ESXi) for compute, NSX for networking, VSAN for storage components in a single stack. VIO is another option for enterprises to run an OpenStack deployment on top of their existing VMware infrastructure.

VIO includes wizard-like installation and configuration with VMware-specific customization and best practices. VIO deploys OpenStack components in active-active clusters, resulting in a highly available configuration.

VIO is comprised of two main building blocks: the VIO Manager and OpenStack components. It is packaged as a single OVA file (Open Virtual Appliance) that contains the Manager server and an Ubuntu Linux virtual machine to be used as the template for the OpenStack components:

- OpenStack controllers - two virtual machines running the following services in an active-active cluster:
    - Dashboard (code-named Horizon)
    - Compute services (Nova API, scheduler and VNC)
    - Identity Service (Keystone)
    - Orchestration Service (Heat)
    - Image Service (Glance)
    - Block Storage (Cinder)
- Memcached cluster
- RabbitMQ cluster, for messaging services used by all OpenStack services
- Load Balancer virtual machines, an active-active cluster managing the internal and public virtual IP addresses
- Nova Compute machine, running the n-cpu service
- Database cluster - a three node MariaDB Galera cluster that stores the OpenStack metadata
- Object Storage machine, running Swift services
- Optional DHCP nodes (required if NSX is not the Neutron provider)

VIO provides VMware monitoring and troubleshooting tools in the OpenStack environment, including VMware vRealize™ Operations™, vRealize™ Log Insight™, and vRealize™ Business™. Administrators can deploy both traditional and cloud native applications with common and familiar VMware tools.

The standard vendor-neutral OpenStack APIs are exposed to developers, as well as the Horizon dashboard. They can take advantage of stored images and flavors for their projects. VIO is a fully standard OpenStack distribution, integrated with VMware technologies as shown in Figure 4.
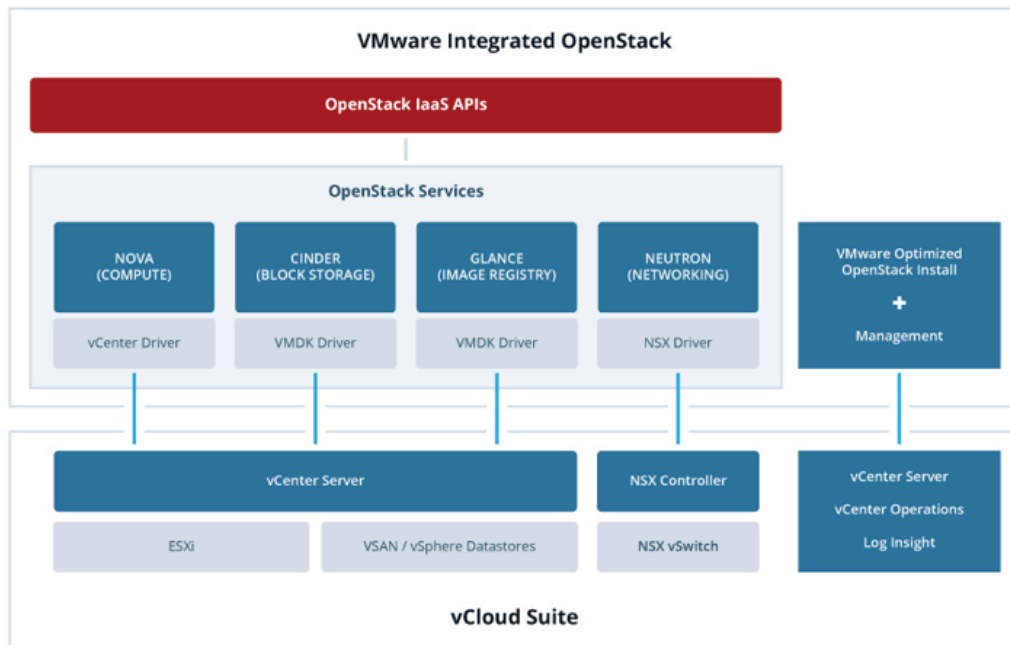
Figure 4: VMware Integrated OpenStack Distribution Architecture

When Adobe® Digital Marketing set out to build their next generation of private cloud, they chose to standardize on OpenStack , to offer self-service, automation and segmentation via a vendor neutral API. They also wanted to scale network, compute and storage independently. They are implementing VIO for several internal SaaS-based applications. These applications require high levels of infrastructure resiliency; hence Adobe was already using and will be retaining VMware HA (High Availability) and DRS (Distributed Resource Scheduling).

VIO enables them to leverage their existing in-house VMware expertise, while well-positioning them for future enhancements. Their key objectives were to quickly offer self-service resource consumption to their product teams, avoiding vendor lock-in and platform maturity.

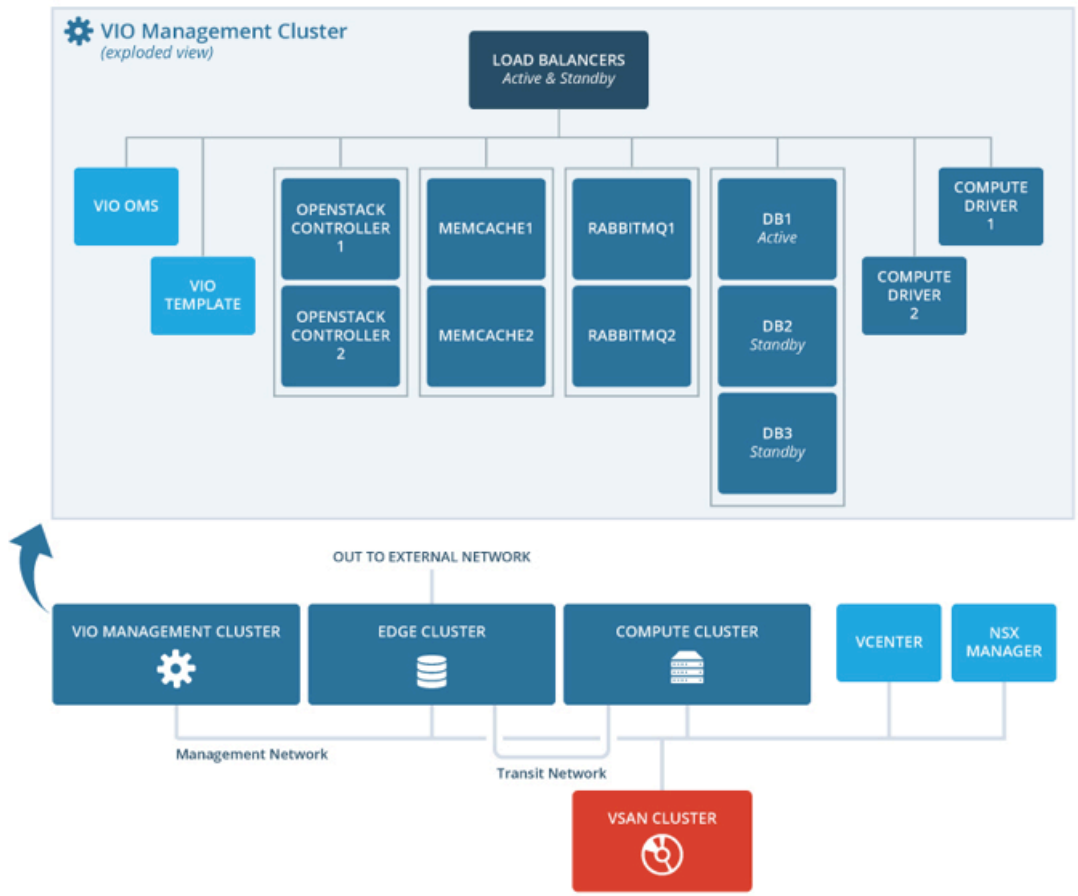Although relatively new, Adobe will be using VIO in production by year-end 2015. Read all about it at **http://openstack.org/assets/pdf-downloads/adobe-dm-case-study.pdf**

Figure 5: Example of VIO User Architecture

Summary

Despite what may seem like overlapping or competitive technology sets, companies can get the best of OpenStack and VMware. Together, they combine the industry leading data center virtualization technologies with the OpenStack cloud to speed new applications to market and help companies compete in a fast moving economy. OpenStack can also help attract and retain skilled developers by offering industry-wide APIs and an ecosystem that cloud application developers love.

Adobe, Intel and the large bank are great examples of how to can combine the agility of OpenStack with vSphere for its compute, network, storage, and management virtualization capabilities.

A basic tour of OpenStack provisioning of compute, storage and network resources with vSphere and NSX is available at **http://www.vmware.com/go/openstacklab**.

Resources to help start an OpenStack cloud, or use an existing one are available at **www.openstack.org/start**. This site serves as a single source for all the information needed to decide about next steps with OpenStack.

Additionally, the OpenStack Marketplace **www.openstack.org/marketplace** offers authorized planning and implementation consultants, public cloud providers supporting ESXi, fully supported OpenStack distributions and more. Contacting one of these providers will help ensure OpenStack success.