5 Things You Need for a True VMware Private Cloud

PLATFORM9

Introduction

IT teams forging their cloud strategies are considering public cloud providers such as Amazon Web Services (AWS) to satisfy growing developer demand for self-service access to infrastructure, automation, and more. But for some organizations and applications, concerns around public cloud security, unpredictable costs, and data gravity/lock-in require IT to consider a private cloud approach that leverages on-premises infrastructure to meet these needs. In most cases, virtualized infrastructure is already in place, raising the question of how to transform these existing resources into a true private cloud.

This white paper:

- Outlines the 5 things you need for a true private cloud.
- Demonstrates how OpenStack delivers a true private cloud environment for VMware vSphere customers.
- Explains how to rapidly create a fully functional private cloud using existing IT staff and infrastructure.

What Makes a True Private Cloud: The Big 5

AWS has made using the public cloud easy and intuitive, providing frictionless on-demand deployment and management of infrastructure services, tailored to meet the needs of applications and the developers who build them. Many organizations desire that public cloud experience for their on-premises infrastructure - in essence, a private cloud. More than just virtualized pooled resources that deliver baseline efficiencies, a true private cloud goes further and includes these 5 key things:

1. Streamlined virtual machine (VM) provisioning
2. Self-service automation
3. Rich application catalog
4. Orchestration using application blueprints
5. AWS-like REST APIs

OpenStack: The Private Cloud Enabler

It's impossible to examine the options for a private cloud without looking at OpenStack. Since its debut in 2010, OpenStack has become the industry's leading private cloud management platform. Interest in OpenStack open source technology

has exceeded alternative private cloud management products such as VMware vRealize Automation (formerly vCloud Automation Center), largely because of its advantages in delivering the 5 key things outlined above.

OpenStack's design is inspired by AWS, providing well-documented REST APIs that enable a self-service, elastic Infrastructure-as-a Service (IaaS) cloud.

As the diagram below shows, OpenStack architecture is both loosely coupled and extensible, making it fundamentally agnostic to the underlying infrastructure because it integrates with various compute, virtualization, network, and storage technologies.

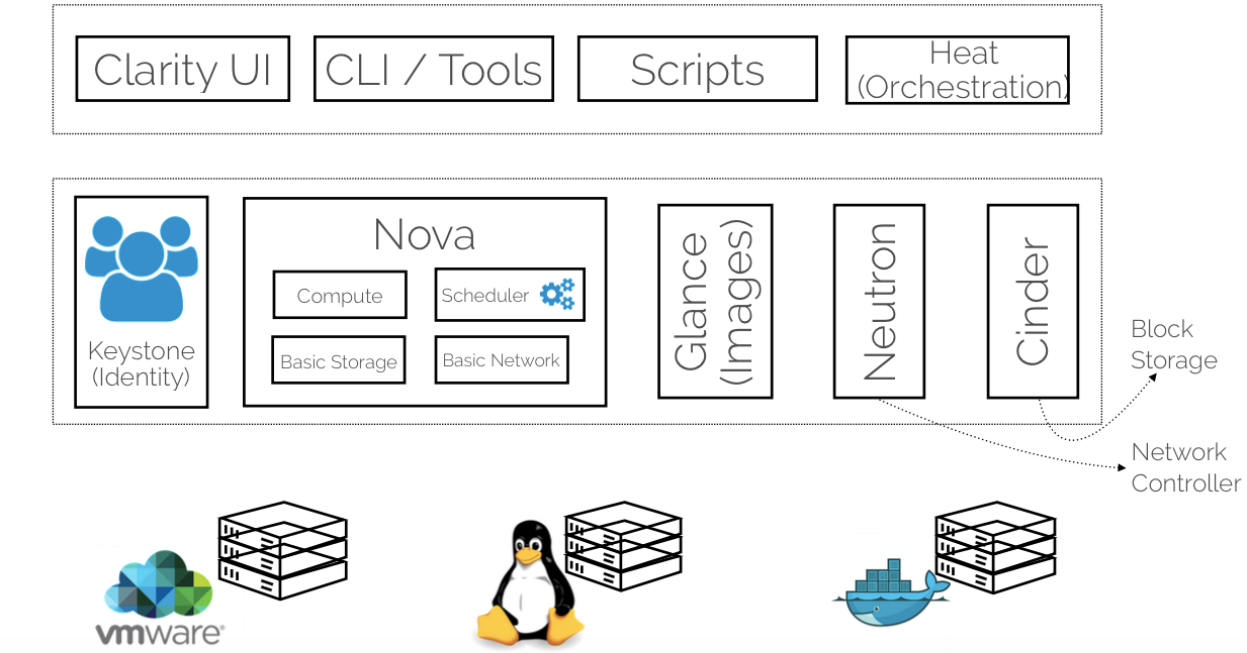

The 5 Key Things: A Closer Look

#1: Streamlined VM Provisioning

A virtualized infrastructure makes it very easy to create new VM instances, which are often kept running for long periods of time. This can lead to organizations ending up with a sprawl of VMs, each running different operating systems, patches, and with inconsistent resource allocation. Ultimately, this sprawl ends up affecting the organization's productivity since it becomes difficult to distinguish between valuable VMs and those just wasting infrastructure capacity.

Hence, the first step to running a private cloud efficiently is to streamline VM provisioning by:

- Centralizing image management to define a curated set of VM templates from which users can create instances (see screenshot).

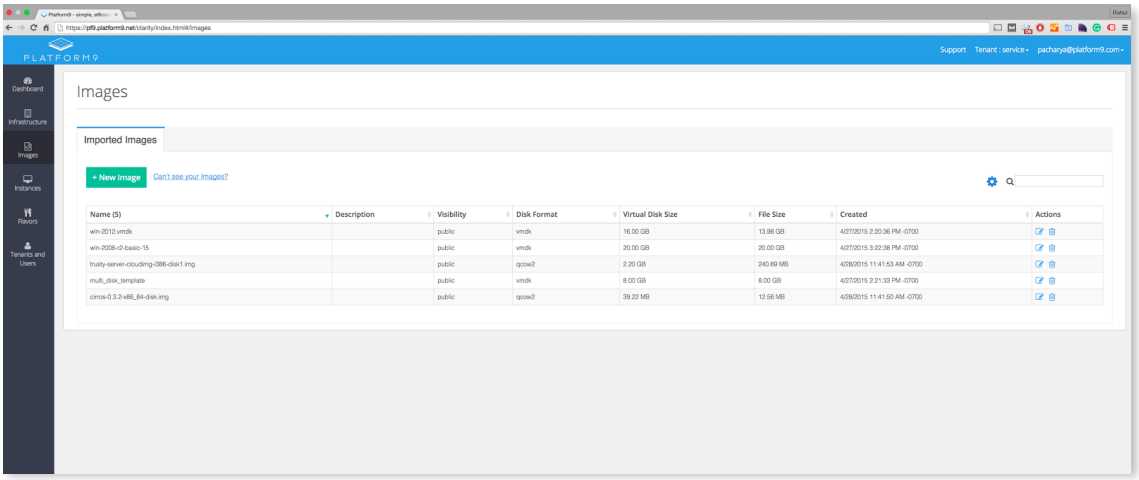

- Reducing the number of unique VM templates by using dynamic machine customization when VMs are created, thereby minimizing the number of unique configurations in the environment.
- Standardizing resource allocations into 'flavors' that capture CPU, memory, disk and other allocation policies for your environment. Flavors apply these predefined resource allocations during VM provisioning (see screenshot).

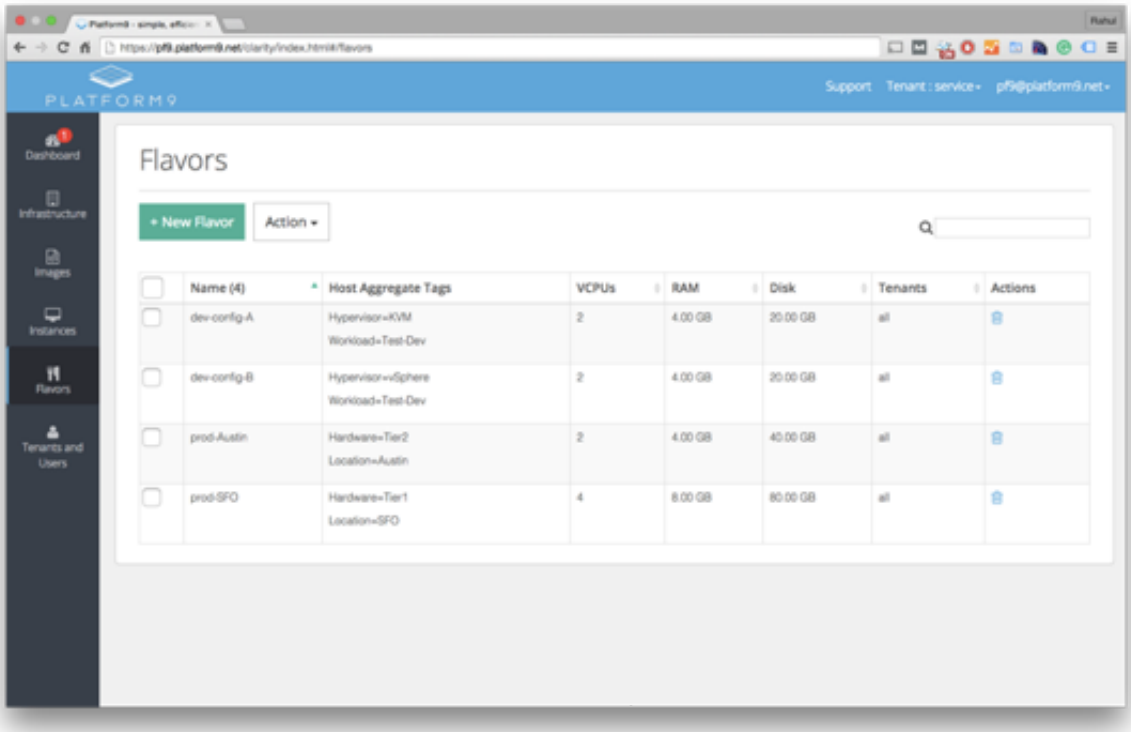

#2: Self-Service Automation

Once VM provisioning is streamlined, the next step is enabling automation, giving users self-service access to infrastructure subject to capacity and policy constraints.

OpenStack makes self-service possible by pooling underlying infrastructure resources and automating VM placement across those pools. The OpenStack scheduler uses resource utilization statistics to place VMs and can automate IP address assignment to newly deployed VM instances, thus speeding up application deployment while requiring fewer manual steps. Administrators can then onboard users and define policies to enable self-service provisioning access. This is far simpler to do in OpenStack than in VMware's vRealize Suite.

For example, consider an IT operations team with one rack of capacity in San Francisco and three racks in Austin. Before OpenStack, the team manually deployed each new VM instance on selected servers - a laborious, error prone task that consumed much of the team's bandwidth.

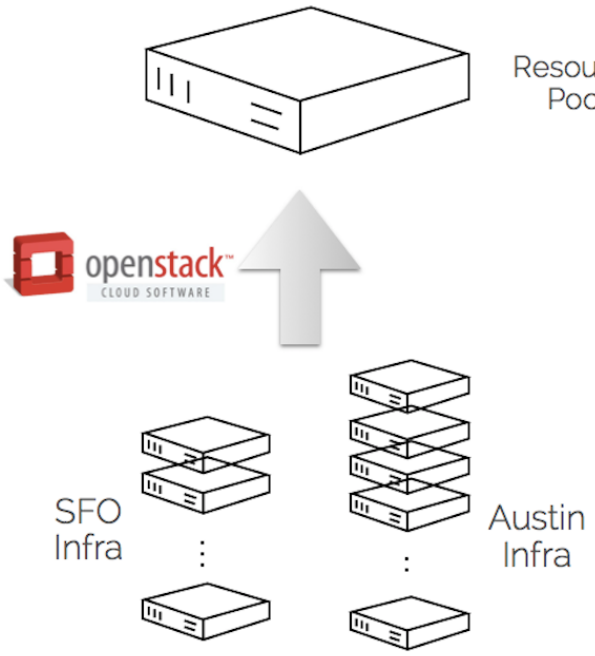

Implementing OpenStack allows the team to manage their combined compute, memory, storage, and networking infrastructure as a single resource pool. The OpenStack scheduler reduces manual effort and human error by automatically determining the best resources to use when deploying each new VM instance.

#3: Rich Application Catalog

Users of the public cloud often start by consuming basic compute, storage and network services, but then go on to require more advanced capabilities such as databases, message queues, load balancers and log processors. Once a private cloud is being used to provision VM based workloads, the natural evolution is to request an out-of-the-box provisioning experience for applications such as these.

OpenStack makes it very easy to leverage a growing, community-maintained ecosystem of applications or to publish applications that are specific to organizations. The OpenStack Murano application catalog delivers a “push-button” experience to provisioning such applications, allowing IT operations to publish a catalog containing Docker-based images, Heat-based image templates or standalone VMs.

#4: Orchestration Using Application Blueprints

Traditionally, virtualization administrators deployed applications by manually deploying individual VMs. Because modern, horizontally scalable applications often involve multiple VMs, deploying and maintaining such applications involves a lot of manual integration after these VMs are deployed.

The cloud computing approach to this problem is to define the infrastructure footprint of such applications using a markup language - such as YAML - and allowing the cloud platform orchestrate the deployment of the individual provisioning steps. This approach has several benefits:

- Applications can be deployed faster - there are fewer manual steps.
- Use of automation means there are fewer chances for error.
- In addition to initial deployment, the cloud platform can automatically orchestrate common lifecycle operations such as scaling the application up or down based on demand.
- Automation blueprints can be built over time and shared across organizations, making it easier to automate increasingly sophisticated applications automatically.

The great news is that VMware users now have access to all of these benefits: when using OpenStack to manage VMware vSphere, the Heat orchestration service enables all of these capabilities, much like the AWS CloudFormations service does for public cloud users.

#5: AWS-Like REST API

All OpenStack services provide a simple, stable, and versioned REST API that works with any language and gives users the freedom to choose from various technologies. This API is a key differentiator between OpenStack and VMware's cloud management platform.

Similar to AWS APIs, OpenStack APIs are easier to program against, more cloud centric, and thus more conducive to providing self-service capabilities to end users than are VMware vSphere APIs. Operations such as creating a VM involve fewer and simpler steps with OpenStack APIs[1] compared to the complex steps required with vSphere APIs[2].

Using an open source API such as the OpenStack API provides many advantages:

- Better ecosystem support: out of the box compatibility with tools such as Vagrant, Terraform, Packer, libcloud (Python), jcloud (Java), Fog (Ruby), Cloud Foundry, Mesos, Puppet, Chef, Salt and Ansible to name a few.
- Ability to leverage new storage and network systems that are OpenStack API compatible.
- Similarity with AWS APIs.
- Avoiding vendor lock-in, a factor that can impact organizations if support for a product/technology is discontinued.

1 http://docs.platform9.com/#a-idcreate-instanceapost-v2tenant_idservers

2 https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.wssdk.apiref.doc_50%2Fvim.vm.ConfigSpec.html

The following table contrasts the OpenStack APIs with proprietary VMware APIs:

	OpenStack API	VMware Proprietary API
Identity Management	• Keystone • Management of users, tenant organizations, and their privileges. Also provides a directory of OpenStack services	• No clear equivalent API
Compute API	• Nova • Exposes the core compute service APIs, including creation of resource tiers, and defining Flavors for standardizing resource allocations	• No clear equivalent API
Image Catalog	• Glance • Exposes a catalog of VM images from which a VM can be provisioned	• No clear equivalent API
Storage APIs	• Cinder (block storage) • Swift (object storage) • The Cinder API can be leveraged by multiple hypervisors and is compatible with a wide range of block storage technologies, including LVM, Ceph, EMC, Hitachi and IBM arrays	• VASA (block storage) is powerful but only compatible with VMware vSphere
Network Control Plane	• Neutron • Multiple plug-ins are available for working with Open vSwitch, Cisco ACI, VMware NSX, Juniper Contrail, PLUMgrid, Linux Bridges, and many others	• Proprietary integration with VMware NSX only
Application Blueprints	• Heat (based on Cloud Formations) • https://github.com/openstack/heat-templates/tree/master/hot	• vApp Director
Application Catalog	• Murano • Enables publishing community built and proprietary application stacks with an easy to use catalog	• No clear equivalent API
Events and Alarms	• Ceilometer • Supports meters and alarms for each of the OpenStack services	• vRealize Operations Manager

Recommended Approach

A recommended approach for organizations that use VMware vSphere is to continue using it as a hypervisor while adopting OpenStack as the management platform. This combination provides organizations with the advantage of using a known, familiar, proven hypervisor while at the same time gaining the advantages available by leveraging OpenStack APIs.

Unfortunately, few IT teams have experienced the full power of OpenStack because of several key challenges:

- As an open source technology, OpenStack does not give users a clear support option when deploying or operating their private cloud.

- OpenStack is complex, requiring specialized skills and knowledge to stand it up and keep it running. Most IT teams do not have, nor can they afford to build, special teams just to manage OpenStack.

For both of these reasons, many OpenStack projects stall, and its potential for powering private clouds remains unrealized.

Platform9's cloud management-as-a-service eliminates these OpenStack challenges, enabling IT to accelerate deployment and simplify the operation of OpenStack-powered private clouds. With Platform9, you can deploy an OpenStack private cloud in minutes to manage new and existing vSphere infrastructures. Platform9 then manages and maintains the control plane while you and your internal customers enjoy an AWS-like private cloud experience as well as the other OpenStack benefits highlighted in this white paper.

To learn more about how Platform9 can simplify and accelerate your private cloud initiatives, contact us at customer-success@platform9.com.

About Platform9

Platform9 transforms an organization's existing servers into an AWS-like agile and efficient self-service private cloud at any scale within minutes while leveraging the latest open source innovations. Powered by OpenStack, Platform9 is the first 100% cloud-managed platform for KVM, VMware vSphere, and Docker. Founded in 2013 by a team of early VMware engineers, Platform9 is situated in Silicon Valley.

Additional Resources

You may also find the following helpful as you evaluate Platform9:

- Blog: Q&A with GE's Cody Hill on the State of the Enterprise Private Cloud
  - http://blog.platform9.com/qa-with-ge-on-the-state-of-the-enterprise-private-cloud
- White Paper: Platform9 for vSphere vs. VMware Integrated OpenStack: 7 Key Differences
  - http://www.platform9.com/resources/whitepaper/Platform9-vs-VMware-Integrated-OpenStack.html
- White Paper: Self-Service and Automation Using OpenStack for VMware vSphere
  - http://platform9.com/resources/whitepaper/Self-Service-Automation-OpenStack-For-VMware.html