

# Getting to Know Apache CloudStack

Joe Brockmeier

PPMC Member Apache CloudStack

[jzb@apache.org](mailto:jzb@apache.org)

# Five Six Characteristics of Clouds

- On-Demand Self-Service
- Broad Network Access
- Resource Pooling
- Rapid Elasticity
- Measured Service
- API Access

# Cloud Service Models



## USER CLOUD a.k.a. SOFTWARE AS A SERVICE

Single application, multi-tenancy, network-based, one-to-many delivery of applications, all users have same access to features.

*Examples: Salesforce.com, Google Docs, Red Hat Network/RHEL*



## DEVELOPMENT CLOUD a.k.a. PLATFORM-AS-A-SERVICE

Application developer model, Application deployed to an elastic service that autoscales, low administrative overhead. No concept of virtual machines or operating system. Code it and deploy it.

*Examples: VMware CloudFoundry, Google AppEngine, Windows Azure, Rackspace Sites, Red Hat OpenShift, Active State Stackato, Appfog*




## SYSTEMS CLOUD a.k.a. INFRASTRUCTURE-AS-A-SERVICE

Servers and storage are made available in a scalable way over a network.

*Examples: EC2, Rackspace CloudFiles, OpenStack, CloudStack, Eucalyptus, OpenNebula*

# Deployment Models



# Why Open Source?

- User-driven solutions to real problems
- Lower barrier to participation
- Larger user base of users helping users
- Aggressive release cycles stay current with the state-of-the-art
- Choice of infrastructure
- Control of data and security
- Avoid lock-in
- *Open data, open standards, open APIs*

History and more...

# UNDERSTANDING CLOUDSTACK

# CloudStack History

- CloudStack began as VM0ps (2008) later became Cloud.com.
- Cloud.com released as CloudStack under GPLv3 (May 2010)
- Cloud.com acquired by Citrix (July 2011)
- Entirely open source (August 2011)
- Relicensed & proposed to Apache (April 3, 2012)
- Accepted as incubating project (April 16, 2012)
- First release (4.0.0-incubating) (November 6, 2012)
- First minor release (4.0.1-incubating) (February 12, 2013)

## Why Apache Software Foundation?



- Best governance
- 15+ years, 100+ projects
- 2500+ Developers



# What is Apache CloudStack?

Apache CloudStack is an open source IaaS that supports multiple hypervisors, high availability, complex networking, firewall, load balancer, and VPN configurations in a multi-tenant environment.

# Specifically...

CloudStack is a set of applications that:

- Provides separation between tenants
- Handle allocating compute resources (inc. custom allocators)
- Let users provision compute resources
- Manages High Availability
- Massively scalable (thousands of nodes)
- Resource usage accounting



The CloudStack UI

# CloudStack UI

- Reference Implementation of the API
- Built with HTML 4, CSS, and jQuery
- Uses Java Server Pages for localization
- Three types of customizations
  - Minor (logo, minor CSS)
  - Major (adding/changing tabs, additional links)
  - Rewrite using API

Dashboard

Instances

Storage

Network

Templates

Accounts

Domains

Events

Infrastructure

Projects

Global Settings

Configuration

General Alerts

View all

- System Alert**  
VM (name: i-5-1683-VM, id: 1683) stopped unexpecte...
- System Alert**  
Unable to restart r-1395-VM which was running on h...
- System Alert**  
Unable to restart r-1410-VM which was running on h...
- System Alert**  
Unable to restart i-31-1424-VM which was running o...
- System Alert**

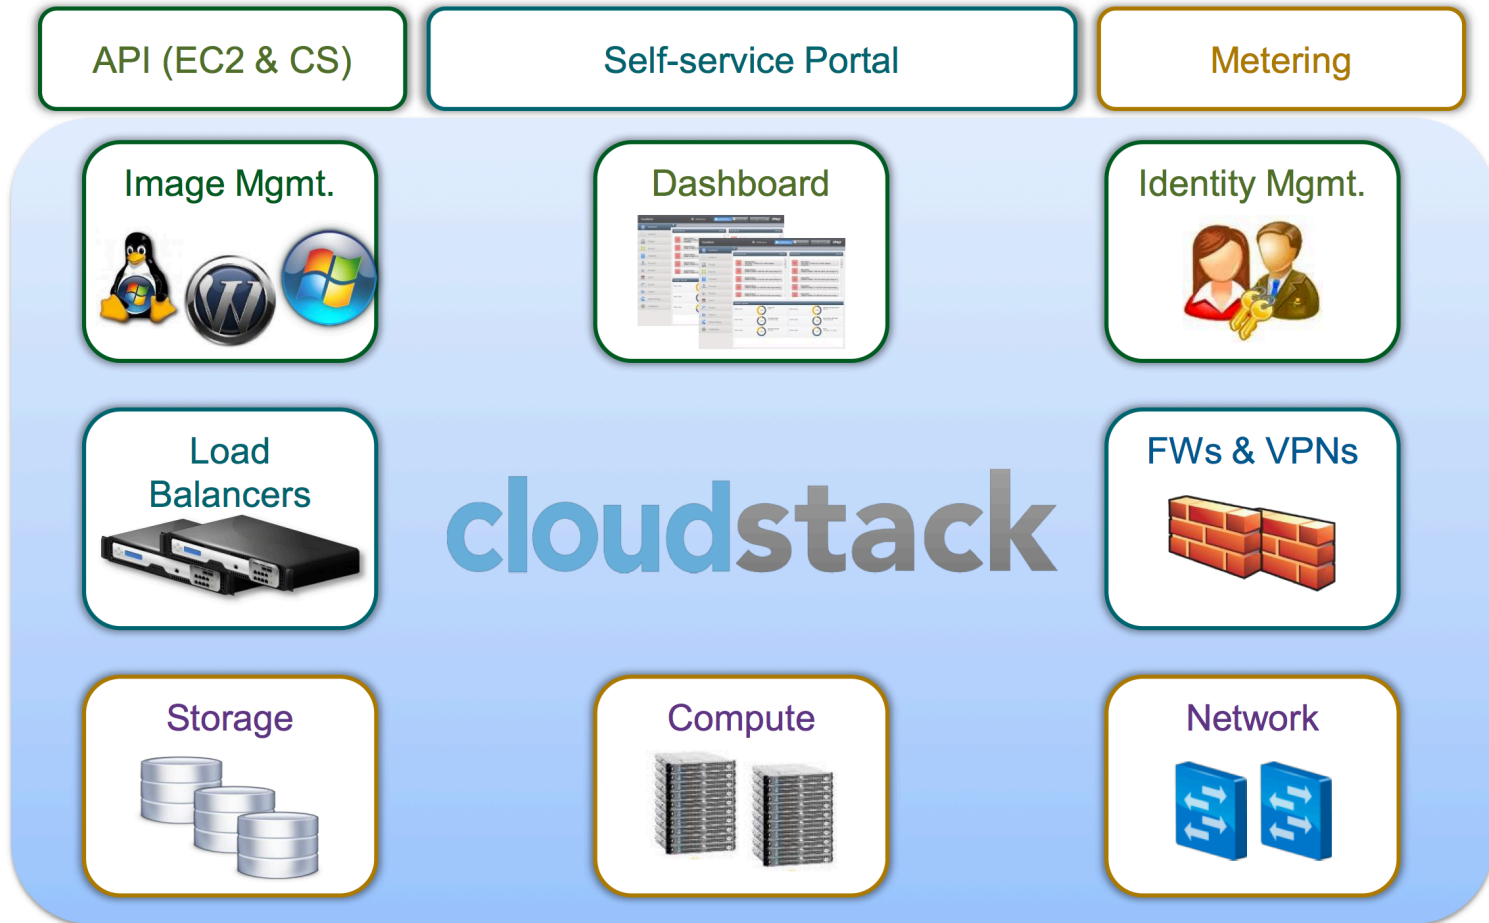
Host Alerts

- Host Alert**  
VM (name: i-5-1683-VM, id: 1683) stopped unexpecte...
- Host Alert**  
Unable to restart r-1395-VM which was running on h...
- Host Alert**  
Unable to restart r-1410-VM which was running on h...
- Host Alert**  
Unable to restart i-31-1424-VM which was running o...
- Host Alert**

System Capacity


Zone: zone		Public IP 2 / 4	Zone: zone		Reserved System IP 40 / 10
Zone: zone		Storage Used 15 GB / 50 GB	Zone: zone		Secondary Storage 15 GB / 50 GB
Zone: zone		Domain Router 100 / 30	Zone: zone		Host 100 gHz / 15.5 gHz

# Or Visually...





# Architectural Overview (Basic)



Conceptual view of a basic deployment

# Hypervisor Support

- KVM
- XenServer
- Xen Cloud Platform
- VMware via vCenter
- Bare Metal via IPMI



# Zone

- Generally used to designate a specific geographic location
- Shares secondary storage across the entire zone
- Has a single network model for the entire zone

# Pod

- General practice – used to designate a rack of machines or row of racks
- Shares a guest network

# Cluster

- Usually a maximum of 8-15 machines per cluster
- Homogeneity is required
  - Same hypervisor (incl. same version)
  - Same CPUs
  - Same networking
- Primary storage is cluster specific

# Secondary Storage

- Stores templates, ISOs, and snapshots
- Historically NFS – added the option of object storage recently
  - Includes Swift, GlusterFS, Ceph and others (in various states of production readiness)
- Managed by Secondary Storage VM –
  - Manages moving templates and snapshots from/to primary storage, aging out snapshots, etc.

# Primary Storage

- Supports NFS, iSCSI, and Clustered Logical Volume Manager (CLVM)
- Requires a shared mountpoint that the hypervisors can mount and write to
- Can use local storage *but* that means no HA, live migration, etc.

# Management Server

- UI/API bits are stateless (state is stored in a MySQL database)
- All UI functionality is available as an API call
- Restful API
  - Unauthenticated API on 8096 for localhost (disabled by default)
  - Authenticated on port 8080
  - Responses in XML or JSON

# CloudStack Allocation

- How are VMs placed, storage allocated, etc.?
- CloudStack has several defaults
  - First fit
  - Fill first
  - Disperse
- Don't like those? Create your own!
- Allows over-provisioning
- OS Preference

# High Availability

- RFMTTR – “really fast mean time to recovery.”
- CloudStack is *not* (alone) a magical HA solution.
- Watches HA-enabled VMs to ensure they’re up, and that the hypervisor it’s on is up. Will restart on another if the hypervisor goes down.
- Redundant router.



# Load Balancing

- Uses HA Proxy
- CloudStack supports load-balancing for distributing traffic to its instances
- Choose between round-robin, source, or least connections
- Choose stickiness policy (source, lbcookie, appcookie)

# Snapshots

- CloudStack allows you to take snapshots manually or set up recurring snapshots.
- Snapshots can be managed automatically (keep N number) and manually (delete snapshot manually).
- Can be turned into templates or volumes to be used by other instances.

# CloudStack Networking

- CloudStack manages
  - DHCP
  - VLAN allocation
  - Firewall
  - NAT/Port forwarding
  - Routing
  - VPN
  - Load Balancing
- CloudStack can manage physical network hardware (F5-Big IP, NetScaler, Juniper SRX)

# CloudStack Networking Types

- CloudStack offers Basic and Advanced Networking
  - Basic: Easy, can only be one physical network
    - Every host, system VM, and guest instance has a unique IP
  - Advanced: Allows multiple physical networks
    - Each account has a public IP, assigned to virtual router, guest IP range (e.g. 10.0.1.10/24), and VLAN ID for the isolated guest network
    - Guests communicate via their own dedicated VLAN

# CloudStack Networks

- Management Network: Used by hypervisors and management server to communicate
- Private Network: Default network for system VMs (virtual router, secondary storage VM, console proxy VM)
- Public Network: Public-facing (e.g. the Internet)
- Guest Network: Network VMs are provisioned on
- Link-local Network: Network used for communication between hypervisor and system VMs. (RFC 3927)

# Security Groups

- Traditional isolation has been via VLAN
- VLANs provide isolation, but at the cost of scaling
  - Standard limit is 4,096 VLANs
  - Hardware that supports upper limit is expensive
  - What happens with 4,097?
- Amazon and others use Layer 3 isolation (Security Groups)

# Security Groups

- Assumes a quasi-trusted Layer 2 network
- Filtering/isolation happens at the bridge level (think ebtables)
- Deny by default

# Accounts, Domains, and Projects

- CloudStack has a top-level domain called ROOT
- You can create sub-domains
- You can create 3 types of accounts, admins, domain-admins, or users
- Projects can be used to hold resources for time-limited projects
- Supports LDAP integration
- CloudStack's account system is very simple – don't make it more complicated!



# Usage Accounting

- Provides stats that can be used for billing (but is not a billing solution)
- Usage stats show: VM count, CPU usage, disk allocation and use, network usage over time.
- Integration howtos for importing to Excel, Ubersmith, Amysta, and Cloud Portal.

# Trying CloudStack

# DevCloud

- DevCloud is a VirtualBox image w/nested virt.

- Grab the DevCloud image from:

<http://wiki.cloudstack.org/display/COMM/DevCloud>

- Log in via the GUI:

<http://localhost:8080/client>

- SSH to DevCloud:

```
ssh -p 2222 root@localhost
```

Username: admin                      Password: password

- DevCloud KVM – effort to run DevCloud on a KVM host

# CloudStack Runbook

- Minimal install (1 server)
- Written by David Nalley w/fixes from the community
- <http://people.apache.org/~ke4qqq/runbook/>
- Focuses on CentOS w/KVM – other runbooks in process for Ubuntu w/Xen or KVM.

# Contact Me



Email: [jzb@apache.org](mailto:jzb@apache.org)



Phone: +1 314-669-4JZB  
(4592)



Blog: <http://dissociatedpress.net/>



Twitter: [@jzb](https://twitter.com/jzb) and [@cloudstack](https://twitter.com/cloudstack)