

It Takes an (Open Source) Village to Build a Cloud



Mark R. Hinkle
Senior Director, Cloud Computing Community
Citrix Systems Inc.

%whoami

- Responsible for management of CloudStack & Xen Open Source Cloud Communities
- Joined Citrix via Cloud.com acquisition July 2011
- Managed Zenoss Open Source project to 100,000 users, 1.5 million downloads
- Former Linux Desktop Advocate (Zealot?)
- Former LinuxWorld Magazine Editor-in-Chief
- Open Management Consortium organizer
- Author - "Windows to Linux Business Desktop Migration" – Thomson
- NetDirector Project - Open Source Configuration Management Project
- Sometimes Author and Blogger at SocializedSoftware.com/ NetworkWorld

The Citrix logo features the word "CITRIX" in a bold, black, sans-serif font. A red dot is positioned above the letter "I", and another red dot is positioned below the letter "X". A registered trademark symbol (®) is located to the upper right of the word.The Xen logo consists of the word "Xen" in a bold, black, sans-serif font. The letter "X" is stylized with a circular outline that passes through its center. A registered trademark symbol (®) is located to the upper right of the word.The cloud.com logo features the word "cloud" in a blue, lowercase, sans-serif font, followed by ".com" in a smaller, grey, lowercase, sans-serif font. A small blue cloud icon is positioned above the letter "d" in "cloud".The Zenoss logo is a black rectangular box with rounded corners. Inside, the word "Zenoss" is written in a white, sans-serif font. Below it, the text "Open Source IT Monitoring" is written in a smaller, white, sans-serif font.The LinuxWorld Magazine logo features the word "LINUXWORLD" in a blue, serif font. A small blue icon of a person is positioned above the letter "L". Below the main text, the word "MAGAZINE" is written in a smaller, blue, sans-serif font.The NetworkWorld logo features the word "NETWORKWORLD" in a bold, blue, sans-serif font. A registered trademark symbol (®) is located to the upper right of the word.The Socialized Software logo features a blue icon of a person with arms raised, followed by the text "Socialized Software" in a blue, sans-serif font.

Quick Cloud Computing Overview: *Or* the Obligatory “What is the Cloud?” Slides

Five Characteristics of Clouds

1. On-Demand Self-Service
2. Broad Network Access
3. Resource Pooling
4. Rapid Elasticity
5. Measured Service

Cloud Computing 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: Google AppEngine, Windows Azure, Rackspace Site, Red Hat Makara




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

Public, Private & Hybrid Clouds



Cloud Still Requires Architectural Design

- Cloud Computing isn't a *magical solution* apps need to be able to *scale out*
- Design your architecture with *the end in mind*
- Make your infrastructure easily *replicable*



Building Clouds with Open Source Software

Private Cloud Architecture



Why Open Source?

- User-Driven Solutions to Real Problems
- Lower barrier to participation
- Larger user base, users helping users
- Aggressive release cycles stay current with the state-of-the-art
- Open data, Open standards, Open APIs



Open Virtual Machine Formats

Open Virtualization Format (OVF) is an open standard for packaging and distributing virtual appliances or more generally software to be run in virtual machines.

Formats for hypervisors/cloud technologies:

- Amazon - AMI
- KVM – QCOW2
- VMware – VMDK
- Xen – IMG
- VHD – Virtual Hard Disk - Hyper-V

Sourcing Open Source Software VMs and Cloud Appliances

Tool/Project	What you can do with them
Bitnami	BitNami provides free, ready to run environments for your favorite open source web applications and frameworks, including Drupal, Joomla!, Wordpress, PHP, Rails, Django and many more.
Boxgrinder	BoxGrinder is a set of projects that help you grind out appliances for multiple virtualization and Cloud providers
SUSE Studio	SUSE Studio supports building and deploying directly to cloud services such as Amazon EC2.
UShareSoft	Create cloud server templates on any OS in minutes. Visually design templates then generate to any image format (hypervisor and physical).







Hypervisors

Open Source

- Xen, Xen Cloud Platform (XCP)
- KVM – Kernel-based Virtualization
- VirtualBox* - Oracle supported Virtualization Solutions
- OpenVZ* - Container-based, Similar to Solaris Containers or BSD Zones
- LXC – User Space *chrooted* installs

Proprietary

- VMware
- Citrix XenServer
- Microsoft Hyper-V
- OracleVM (Based on OS Xen)



OpenVZ



Compute Clouds (IaaS)

	Year Started	License	Virtualization Technologies
CloudStack	2008	Apache	Xenserver, Xen Cloud Platform, KVM, VMware
Eucalyptus	2006	GPL	Xen, KVM, VMware (commercial version)
OpenStack	2010 (Developed by NASA by Anso Labs previously)	Apache	VMware ESX and ESXi, , Xen, Xen Cloud Platform KVM, LXC, QEMU and Virtual Box
OpenNebula	2005	Apache	Xen, KVM, VMware


Numerous companies are building cloud software on OpenStack including Nebula, Piston Inc., CloudScaling

Scale-Up or Scale-Out

Vertical Scaling (Scale-Up) – Allocate additional resources to VMs, requires a reboot, no need for distributed app logic, single-point of OS failure



Horizontal Scaling (Scale-Out) – Application needs logic to work in distributed fashion (e.g. HA-Proxy and Apache, Hadoop)



Cloud Computing Storage

	Description
GlusterFS	Scale Out NAS system aggregating storage over Ethernet or Infiniband
CEPH	Distributed file storage system developed by DreamHost
OpenStack Swift	Long-term object storage system
Sheepdog	Distributed storage for KVM hypervisors
NFS	<i>Old standby, tried and true, not designed for cloud scale or performance</i>

Cloud APIs Aren't Created Equal

Open Source Abstractions

- [jclouds](#)
- [libcloud](#)
- [deltacloud](#)
- [fog](#)



Platform-as-a-Service (PaaS)

	Year Started	Sponsors	Languages/ Frameworks
CloudFoundry	2011	VMware	Spring for Java, Ruby for Rails and Sinatra, node.js, Grails, Scala on Lift and more via partners (e.g. Python, PHP)
OpenShift **	2011	Red Hat	Java, Ruby, PHP, Perl and Python
PHPFog *	2011	Appfog	PHP, NodeJS, Ruby, Python, Java, .NET, MySQL, PostgreSQL
Stackato *		ActiveState	Java, Python, PHP, Ruby, Perl, Node.js, others
WSO2 Stratus	2010	WSO2	Jboss, Java EE6

* Built on Cloud Foundry, some contribution upstream


** Not open source yet

Managing Clouds with Open Source Tools




Automation Unlocks the Potential of the Cloud

Meat Cloud



Cloud Operations



4 Types of Management Tools

Provisioning

Installation of operating systems and other software

Configuration Management

Sets the parameters for servers, can specify installation parameters

Orchestration/Automation

Automate tasks across systems

Monitoring


Records errors and health of IT infrastructure



Management Toolchains

Toolchain (n):

A set of tools where the output of one tool becomes the input of another tool



Open Source Provisioning Tools

	Year Started	License	Installation Targets
Kickstart	?	GPL	Most .dep and RPM based Linux distros
Cobbler (Plus koan for PXE boot of VMs)	2007	GPL	Red Hat, OpenSUSE Fedora, Debian, Ubuntu
Spacewalk	2008	GPL	Fedora, Centos
Crowbar	2011	Apache	(Bare metal provisioning)

Configuration Management Tools

	Year Started	Language	License	Client/Server
Cfengine	1993	C	Apache	Yes
Chef	2009	Ruby	Apache	Chef Solo – No Chef Server - Yes
Puppet	2004	Ruby	GPL	Yes & standalone
Salt	2011	Python	Apache	yes

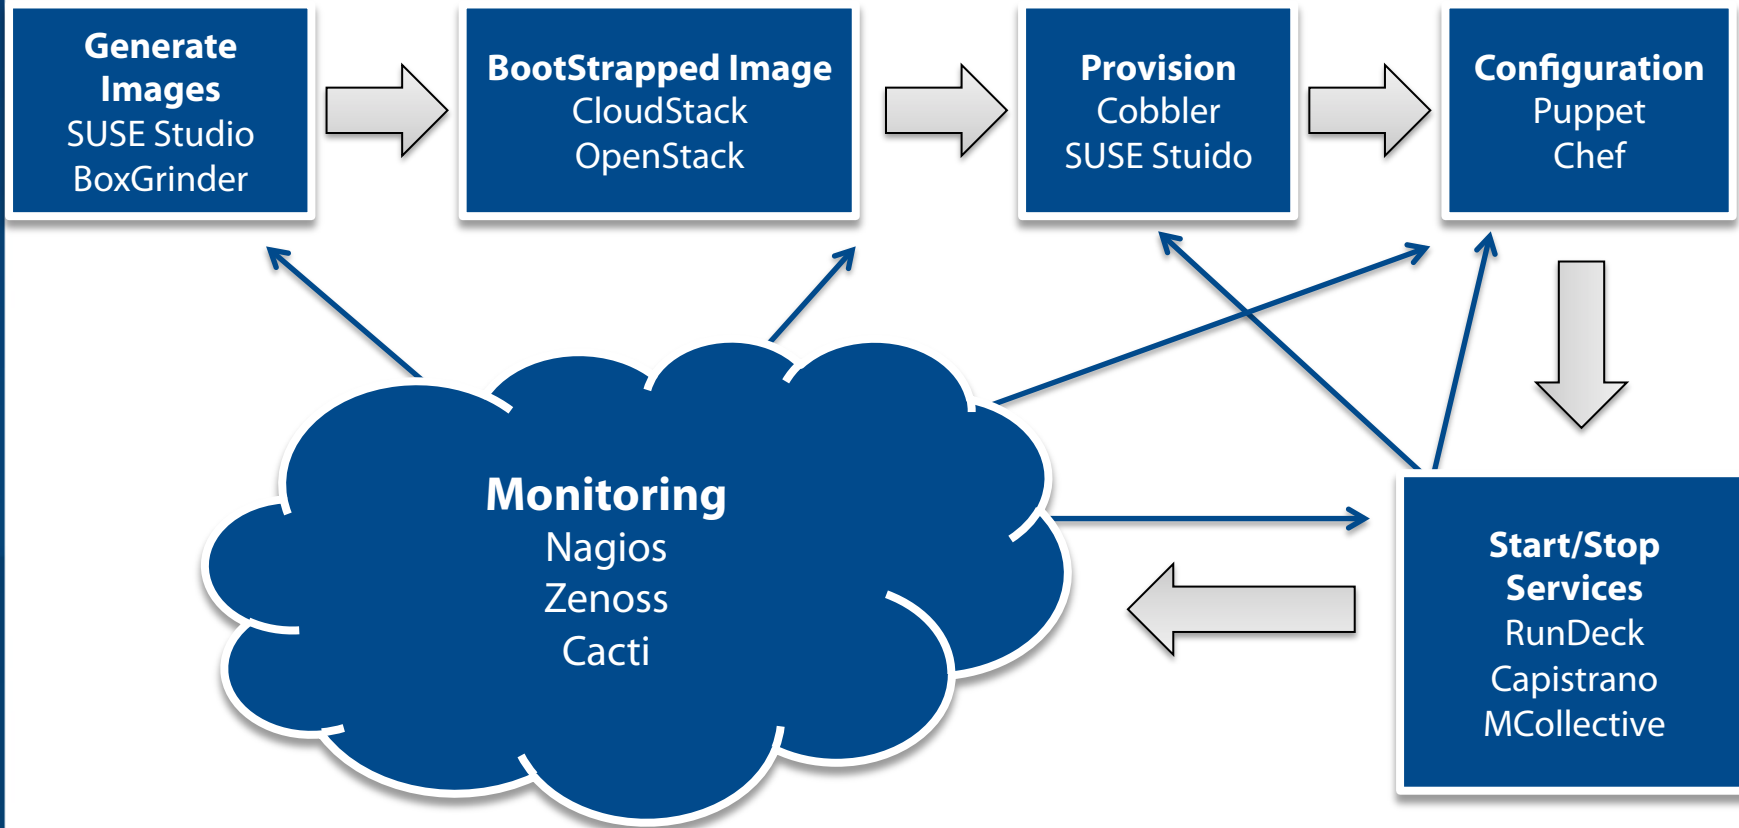
Monitoring Tools

	License	Type of Monitoring	Collection Methods
Cacti / RRDTool	GPL	Performance	SNMP, syslog
Graphite	Apache 2.0	Performance	Agent
Nagios	GPL	Availability	SNMP,TCP, ICMP, IPMI, syslog
Zabbix	GPL	Availability/ Performance and more	SNMP, TCP/ICMP, IPMI, Synthetic Transactions
Zenoss	GPL	Availability, Performance, Event Management	SNMP, ICMP, SSH, syslog, WMI

Automation/Orchestration Tools

	Year Started	Language	License	Client/Server	Support Organization
Capistrano	2006	Ruby	MIT	Yes	None
RunDeck	2010	Java	Apache	Yes	DTO Solutions
Func	2007	Python	GPL	Yes	Fedora Project
MCollective	2009	Ruby	Apache	Yes	PuppetLabs
Salt	2011	Python	Apache	Yes	SaltStack Inc. ?

Conceptual Automated Toolchain



Questions?

Slides Can be Viewed and Downloaded at:

<http://www.slideshare.net/socializedsoftware/>

Contact Me



Professional: mark.hinkle@citrix.com

Personal: mrhinkle@gmail.com



Professional: 919.228.8049



Professional: <http://www.cloudstack.org>

Personal: <http://www.socializedsoftware.com>



Twitter: [@mrhinkle](https://twitter.com/mrhinkle)



Mark R. Hinkle

Director, Cloud Computing Community
Citrix Systems Inc.
Open Source Enthusiast

Additional Resources

- [Devops Toolchains Group](#)
- [DevOps Wikipedia Page](#)
- [Open Cloud Initiative](#)
- [NIST Cloud Computing Platform](#)
- [Open Virtualization Format Specs](#)
- [Clouderati Twitter Account](#)
- [Planet DevOps](#)

It Takes a (Open Source) Village to Build a Cloud

by Mark R. Hinkle is licensed under a Creative Commons Attribution-ShareAlike 3.0 United States License.

