

# Overview: Building Open Source Cloud Computing Environments

LinuxCon Europe 2011  
Prague, Czech Republic

Mark R. Hinkle  
Director, Cloud Computing Community  
Citrix Systems Inc.  
Twitter: [@mrhinkle](https://twitter.com/mrhinkle)  
Email: [mrhinkle@cloudstack.org](mailto:mrhinkle@cloudstack.org)



# Agenda

1. Introduction
2. Quick Cloud Computing Overview
3. Open Source Building Blocks for Cloud Computing
4. Open Source Tools for Cloud Management
5. Questions

# %whoami

- Responsible for Driving Adoption of CloudStack Open Source Cloud Computing Software
- Joined Citrix via Cloud.com acquisition July 2011
- Former manager of Zenoss Open Source project 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](http://SocializedSoftware.com/NetworkWorld)

The Citrix logo features the word "CITRIX" in a bold, black, sans-serif font. Two red dots are positioned above the 'I' and below the 'X'.The cloud.com logo consists of the word "cloud" in a blue, lowercase, sans-serif font, followed by ".com" in a smaller, grey, lowercase, sans-serif font. A blue cloud icon is positioned above the 'd' in "cloud".The Zenoss logo is a black rounded rectangle containing the word "Zenoss" in white, with a bar over the 'o'. Below it, the text "Open Source IT Monitoring" is written in a smaller white font.The LinuxWorld Magazine logo features the word "LINUXWORLD" in a blue, serif font, with a small Linux penguin icon above the 'I'. Below it, 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, with a registered trademark symbol. A red horizontal line is positioned below the text.The Socialized Software logo features a blue icon of three stylized human figures on the left, 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*



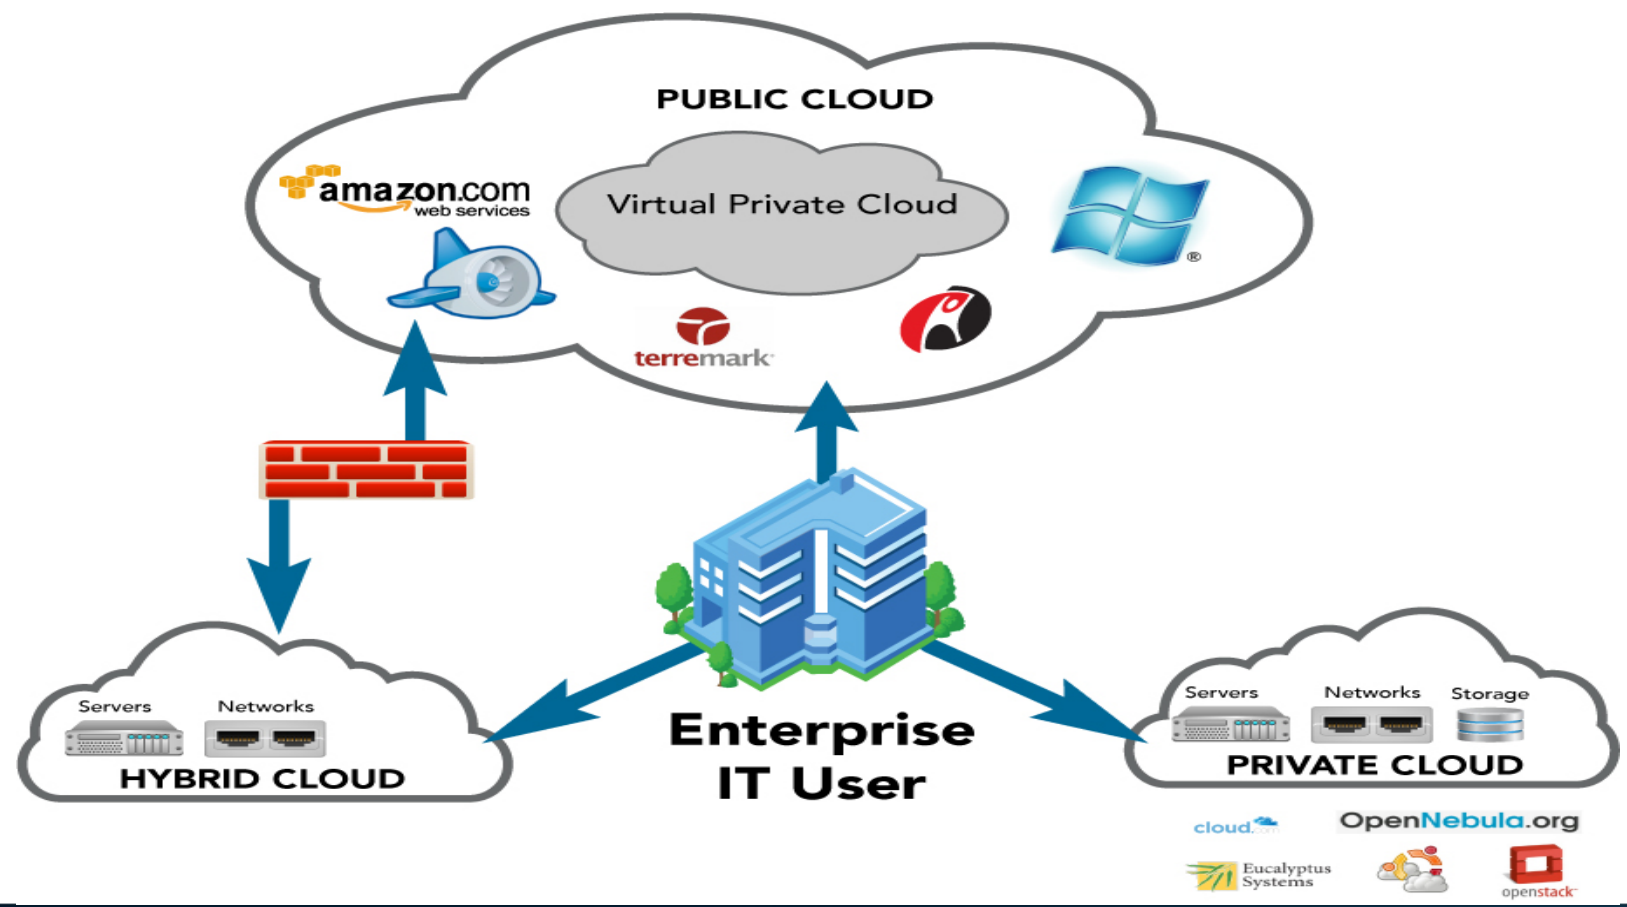
## SYSTEMS CLOUD a.k.a. INFRASTRUCTURE-AS-A-SERVICE

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

*Examples: EC2, Rackspace CloudFiles, OpenStack, CloudStack, Eucalyptus, Ubuntu Enterprise Cloud, OpenNebula*

# Deployment Models

## Public, Private & Hybrid Clouds



# Cloud Still Requires Architectural Design

- Cloud Computing isn't a *magical solution*
- Design your architecture with *the end in mind*
- Make your infrastructure easily *replicable*

# Building Compute Clouds with Open Source Software

**cloudstack**  
open source cloud computing

## PRIVATE OPEN SOURCE CLOUD COMPUTING

	XEN	OPENSTACK
	KVM	CLOUDSTACK
	GLUSTER	COLLECTD
	CEPH	RABBITMQ
ARCHITECT: YOU	JCLOUDS	ZENOSS

# 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 Cloud Initiative (OCI)

Open Cloud Initiative (OCI)

A non-profit advocate of open cloud computing



## Principles of Open Cloud

- *Interoperability* (the ability to exchange and use information) between cloud computing products and services is **required** for unfettered competition between vendors and unrestricted choice for users.
- Users **must** be able to come (*no barriers to entry*) and go (*no barriers to exit*) regardless of who they are (*no discrimination*) and what systems they use (*technological neutrality*).
- Supporting vendors **must** therefore cooperate on standards, implementing those that exist (where applicable) and collaborating via an *open process* to develop those that don't, with a view to competing fairly on quality.

<http://www.opencloudinitiative.org>

# 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

# Sourcing OSS

## VMs and **Cloud Appliances**



# BoxGrinder

# BITNAMI

# Open Source 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



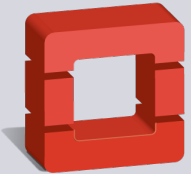
- VMware
- Citrix XenServer
- Microsoft Hyper-V
- OracleVM (Based on OS Xen)



OpenVZ






# Open Source Compute Clouds

	Year Started	License	Hypervisors Supported
 open source cloud computing	2008	GPL	Xenserver, Xen Cloud Platform, KVM, VMware
 Eucalyptus Systems	2006	GPL	Xen, KVM, VMware
 openstack™	2010 (Developed by NASA by Anso Labs previously)	Apache	VMware ESX and ESXi, Microsoft Hyper-V, Xen, KVM and Virtual Box

Other open source compute software include Abiquo, Red Hat's CloudForms and OpenNebula. Numerous companies are building cloud software on OpenStack including Nebula, Piston Inc.

# Open Source Platform-as-a-Service

	Year Started	Sponsors	Hypervisors Supported
 <p>CLOUD FOUNDRY™</p>	2011	VMware	Spring, Rails, Node.js
 <p>SmartOS</p>	2011	Joyent	Node.js
 <p>OPENSHIFT™ PaaS by Red Hat® Cloud</p>	2011	Red Hat	Express – Ruby, Php and Python Flex – Jboss, Java EE6
 <p>WSO2 Stratos</p>	2010	WSO2	Jboss, Java EE6

# Open Source Cloud Computing Storage

- [GlusterFS](#) – Scale Out NAS system aggregating storage over Ethernet or Infiniband
- Ceph – Distributed file storage system developed by DreamHost
- [OpenStack Object Storage \(SWIFT\)](#) – Long-term storage object storage system
- [Sheepdog](#) – Distributed storage for KVM hypervisors
- *NFS – Old standby, tried and true, not designed for cloud scale or performance*

# Cloud APIs Aren't Created Equal

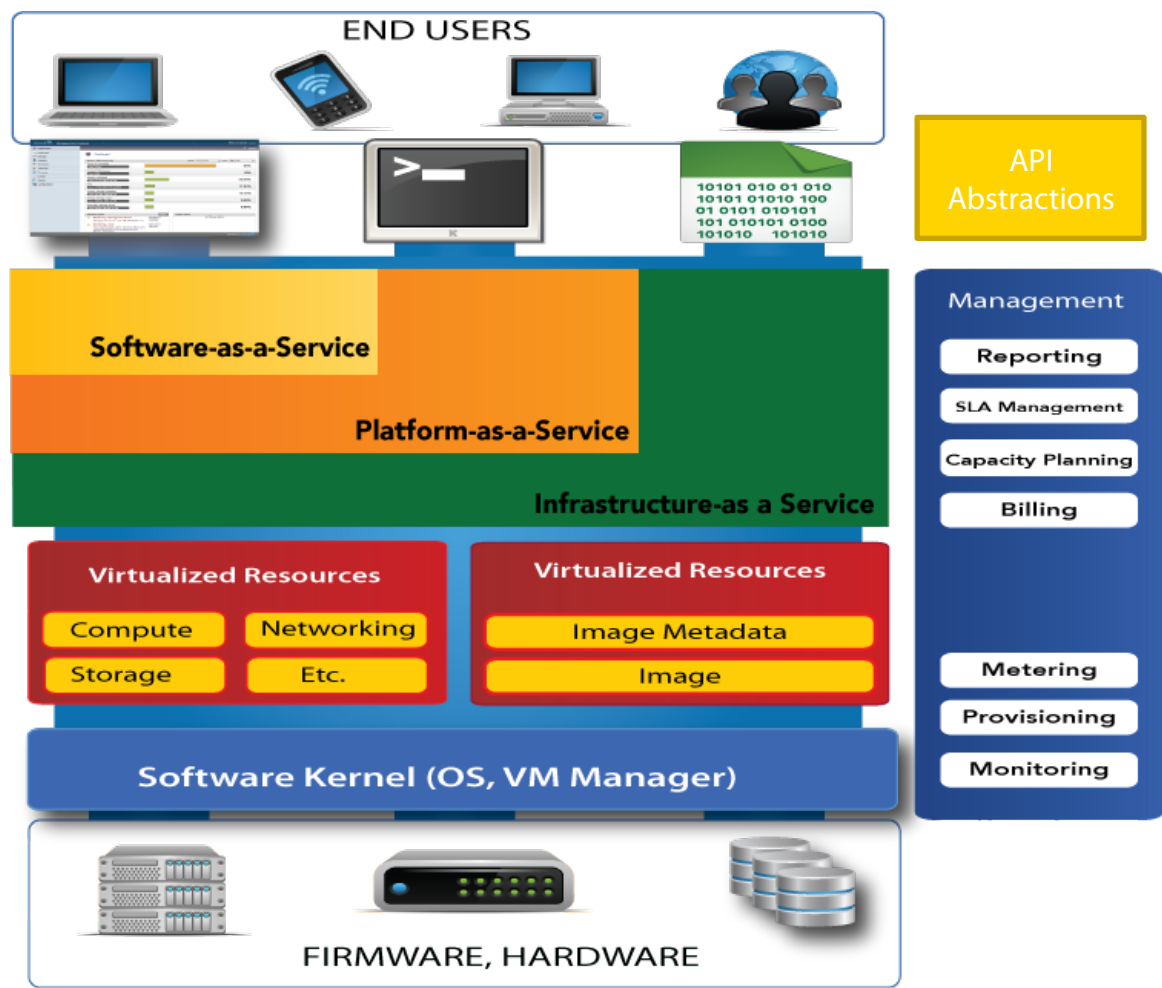
## Open Source Abstractions

18


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



# Private Cloud Architecture




# Managing Clouds with Open Source Tools



# Automation Unlocks the Potential of the Cloud

- MeatCloud, Can't Keep up with Cloud Computing
- Devops & Agile IT Philosophy
- Script Repetitive Tasks
- *Automate, Automate, Automate*



# The Myth of the Nines

Availability %	Downtime per Year	Downtime per Month	Downtime per Week
99.9% (three nines)	8.76 hours	43.2 minutes	10.1 minutes
99.95%	4.38 hours	21.56 minutes	5.04 minutes
99.99% (four nines)	52.6 minutes	4.32 minutes	1.01 minutes
99.999% (five nines)	5.26 minutes	25.9 seconds	6.05 seconds
99.9999% (six nines)	31.5 seconds	2.59 seconds	.0605 seconds

Average polling interval for monitoring - 5 minutes  
 Even superhuman operations people can't be alerted and take action in under 5 minutes.  
 One outage per year could drop service level to three nines or worse.

# 4 Types of Management Tools

```

user@linux:~$ yum install lynx
Gathering header information file(s) from server(s)
Server: Fedora Linux / stable for Red Hat Linux 9 (i386)
Server: Fedora Linux / testing for Red Hat Linux 9 (i386)
Server: Red Hat Linux 9 (i386)
Server: Red Hat Linux 9 (i386) updates
Finding updated packages
Downloading needed headers
getting /var/cache/yum/fedora-stable/headers/leafnode-0-1.9.43-
1.fdr.1.i#90.i386.hdr
getting /var/cache/yum/fedora-stable/headers/libzvt-devel-0-2.0.1-
1.fdr.1.i#90.i386.hdr
getting /var/cache/yum/fedora-stable/headers/mhash-devel-0-0.8.18-
1.fdr.1.i#90.i386.hdr
  
```

## Provisioning


Installation of operating systems and other software

## Configuration Management

Sets the parameters for servers, can specify installation parameters

```

[user@linux:~]$ vim http.conf
#
# Based upon the NCSA server configuration files originally by Rob McCool
#
# This is the main Apache server configuration file.  It contains the
# configuration directives that give the server its instructions.
# See http://www.apache.org/docs/ for detailed information about
# the directives.
#
# Do NOT simply read the instructions in here without understanding
# what they do.  They're here only as hints or reminders.  If you are unsure
# consult the online docs.  You have been warned.
#
# After this file is processed, the server will look for and process
  
```



## Orchestration/Automation

Automate tasks across systems

## Monitoring

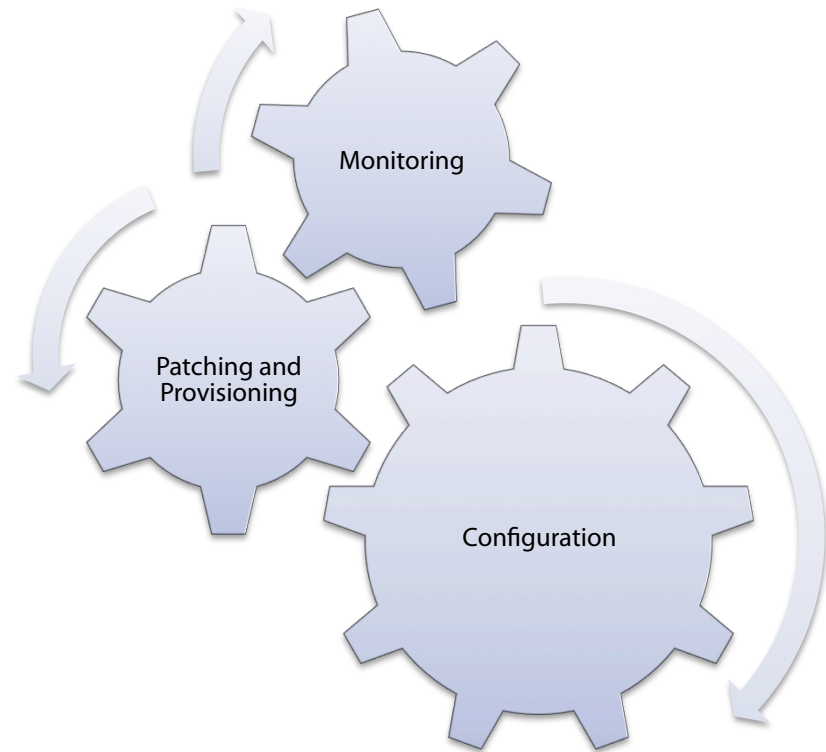
Records errors and health of IT infrastructure



# 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	Language	License	Installation Targets
Kickstart	?	Python	GPL	Most .dep and RPM based Linux distros
<a href="#">Cobbler</a> (Plus koan for PXE boot of VMs)	2007	Python	GPL	Red Hat, OpenSUSE, Fedora, Debian, Ubuntu
<a href="#">Spacewalk</a>	2008	Perl, Python, Java	GPL	Fedora, Centos
<a href="#">Crowbar</a>	2011	Ruby	Apache	(Bare metal provisioning)

# Open Source Configuration Management Tools

	Year Started	Language	License	Client/Server
<a href="#">Cfengine</a>	1993	C	Apache	Yes
<a href="#">Chef</a>	2009	Ruby	Apache	Chef Solo – No Chef Server - Yes
<a href="#">Puppet</a>	2004	Ruby	GPL	yes

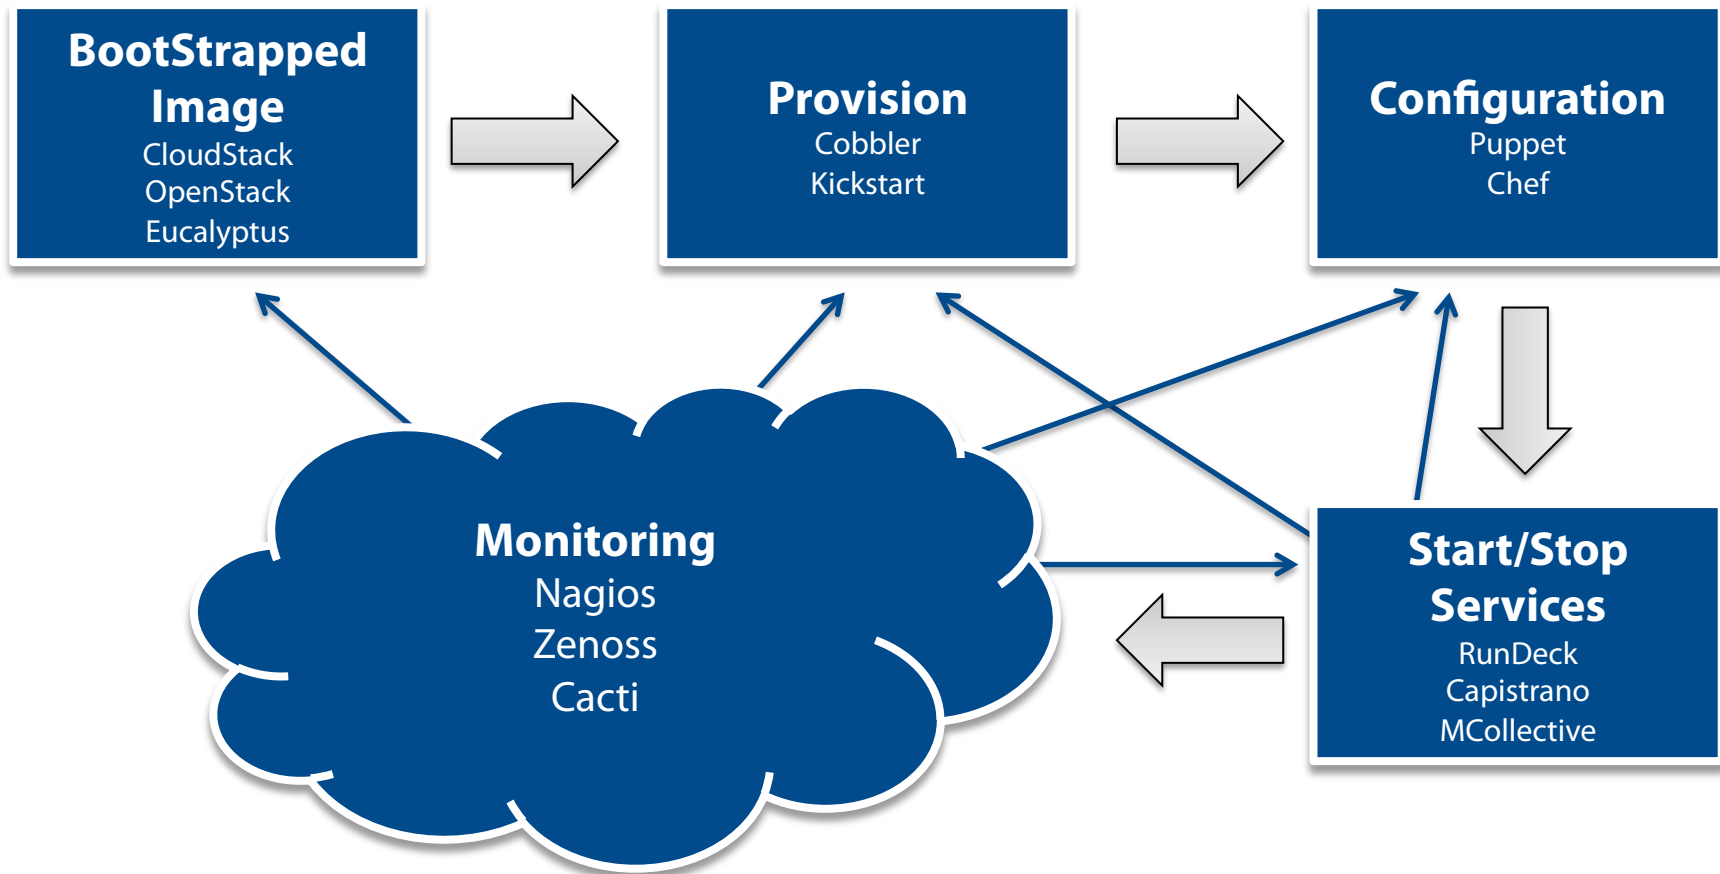
# Open Source Monitoring Tools

	License	Type of Monitoring	Collection Methods
<a href="#">Cacti / RRDTool</a>	GPL	Performance	SNMP, syslog
<a href="#">Nagios</a>	GPL	Availability	SNMP, TCP, ICMP, IPMI, syslog
<a href="#">Zabbix</a>	GPL	Availability/ Performance and more	SNMP, TCP/ICMP, IPMI, Synthetic Transactions
<a href="#">Zenoss</a>	GPL	Availability, Performance, Event Management	SNMP, ICMP, SSH, syslog, WMI

# Open Source Automation/Orchestration Tools

	Year Started	Language	License	Client/Server	Support Organization
<a href="#">Capistrano</a>	2006	Ruby	MIT	Yes	None
<a href="#">RunDeck</a>	2010	Java	Apache	Yes	DTO Solutions
<a href="#">Func</a>	2007	Python	GPL	Yes	Fedora Project
<a href="#">MCollective</a>	2009	Ruby	Apache	Yes	PuppetLabs

# Conceptual Automated Toolchain



# Questions?

**Slides Can be Viewed and Downloaded at:**

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

# Contact Me



Professional: [mrhinkle@cloustack.org](mailto:mrhinkle@cloustack.org)

Personal: [mrhinkle@gmail.com](mailto:mrhinkle@gmail.com)



Professional: 919.228.8049



Professional: <http://www.cloustack.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](#)

# Crash Course in Open Source Cloud Computing

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

