Crash Course in Open Source Cloud Computing

David Nalley
Cloudstack Community Guy
Cloud.com
IRC: ke4qqq on irc.freenode.net
Email: david.nalley@cloud.com

cloud.com
Open Source Cloud Computing Software

%whoami

- NOT Mark Hinkle
- Community guy for Cloudstack
- Fedora Project Board member (along with docs writer, packager, infrastructure sysadmin, and a host of other roles)
- Organizer of the Southeast Linuxfest
- Recovering sysadmin
- Contributor to a lesser degree for other F/LOSS projects (Zenoss, Sahana, Cobbler, Sheepdog)
- Author for a number of magazines and websites. (namely Linux Pro Magazine and Opensource.com)

Slides Can be Viewed and Downloaded at:

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

Agenda

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

Cloud Computing:
Cloud Computing Trends

http://www.dilbert.com/strips/comic/2011-01-07/

Recent Study on **Cloud Computing Preferences**

Top Reasons For Cloud Computing Adoption

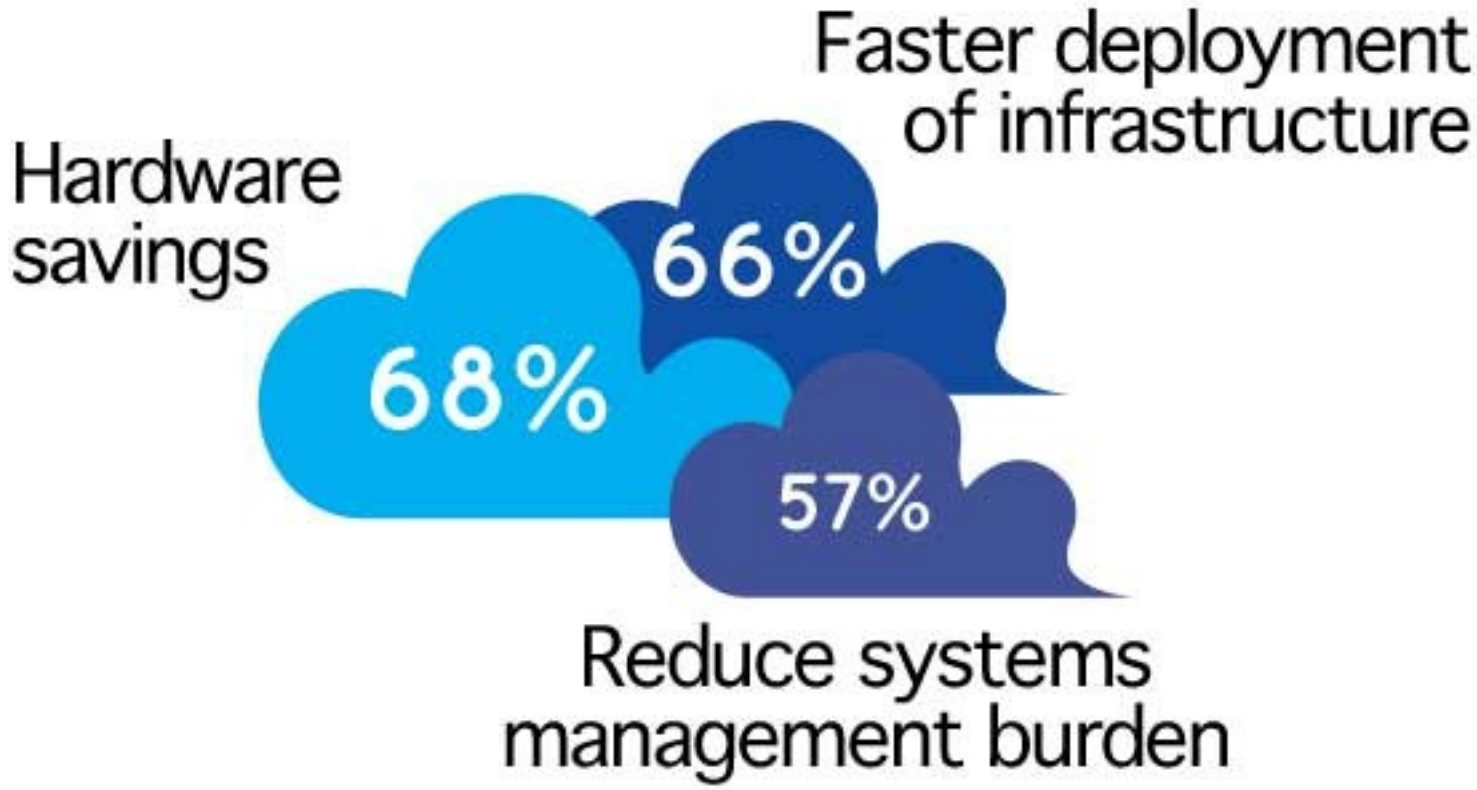

Open Source Usage

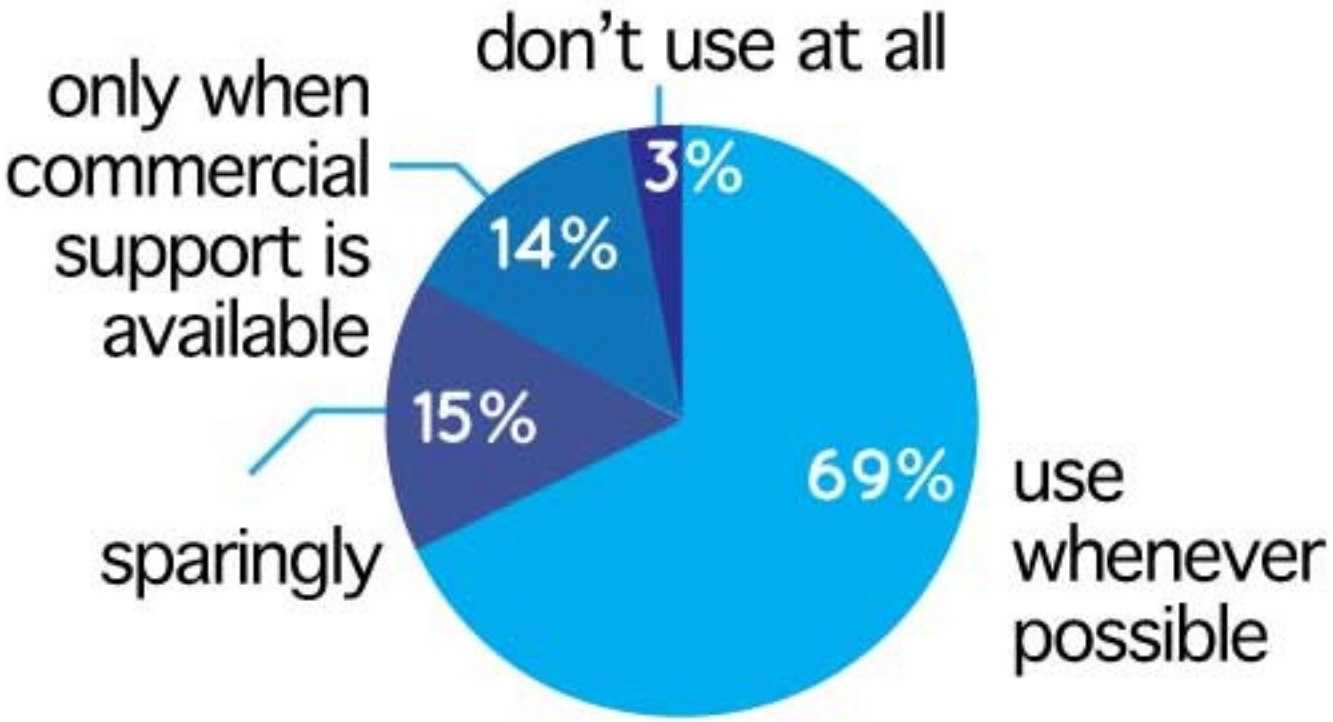

Cloud Service Model Usage

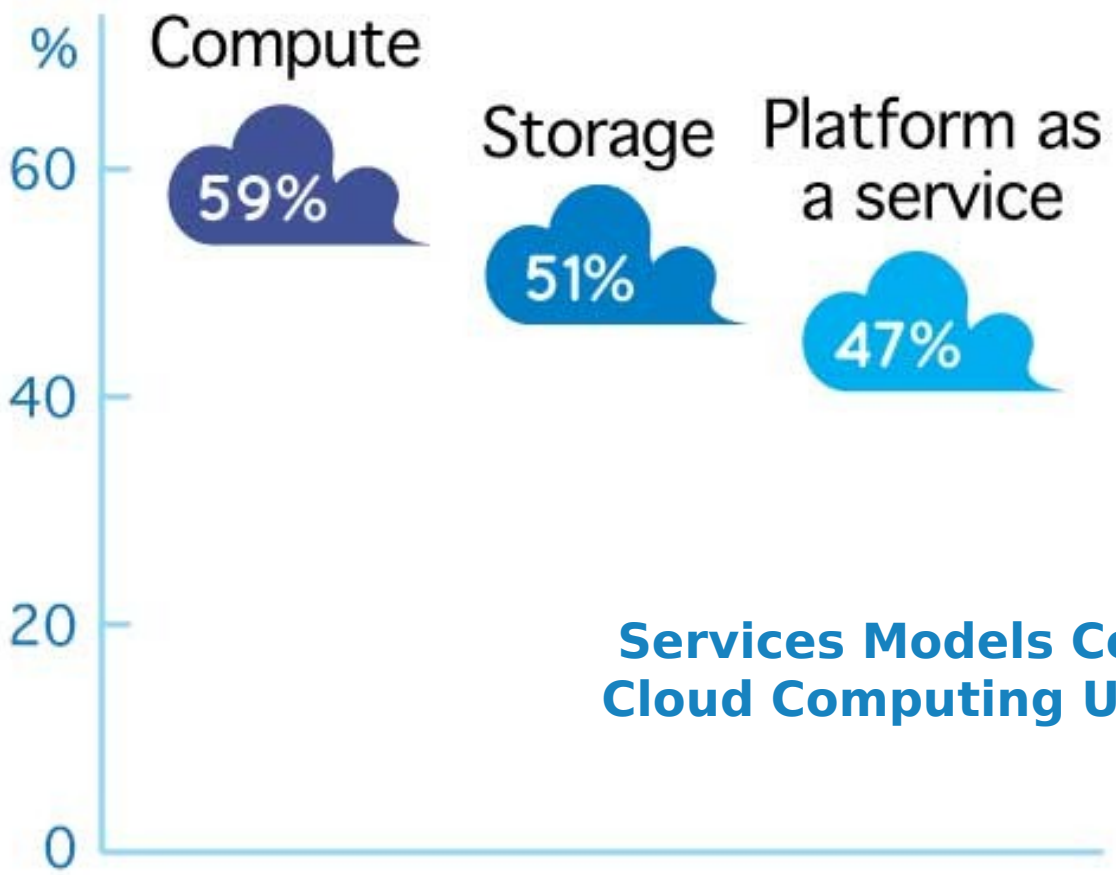

Services Models Consumed by Cloud Computing Users in 2011

Why Open Source?

- Typically User-Driven to solve real problems
- Larger user base, users helping users
- Lower barrier to participation
- Aggressive release cycles – stay current with the state of the art
- Try before you "buy", no Brochure-ware, no Powerpoint software
- Open data, Open standards, Open APIs

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

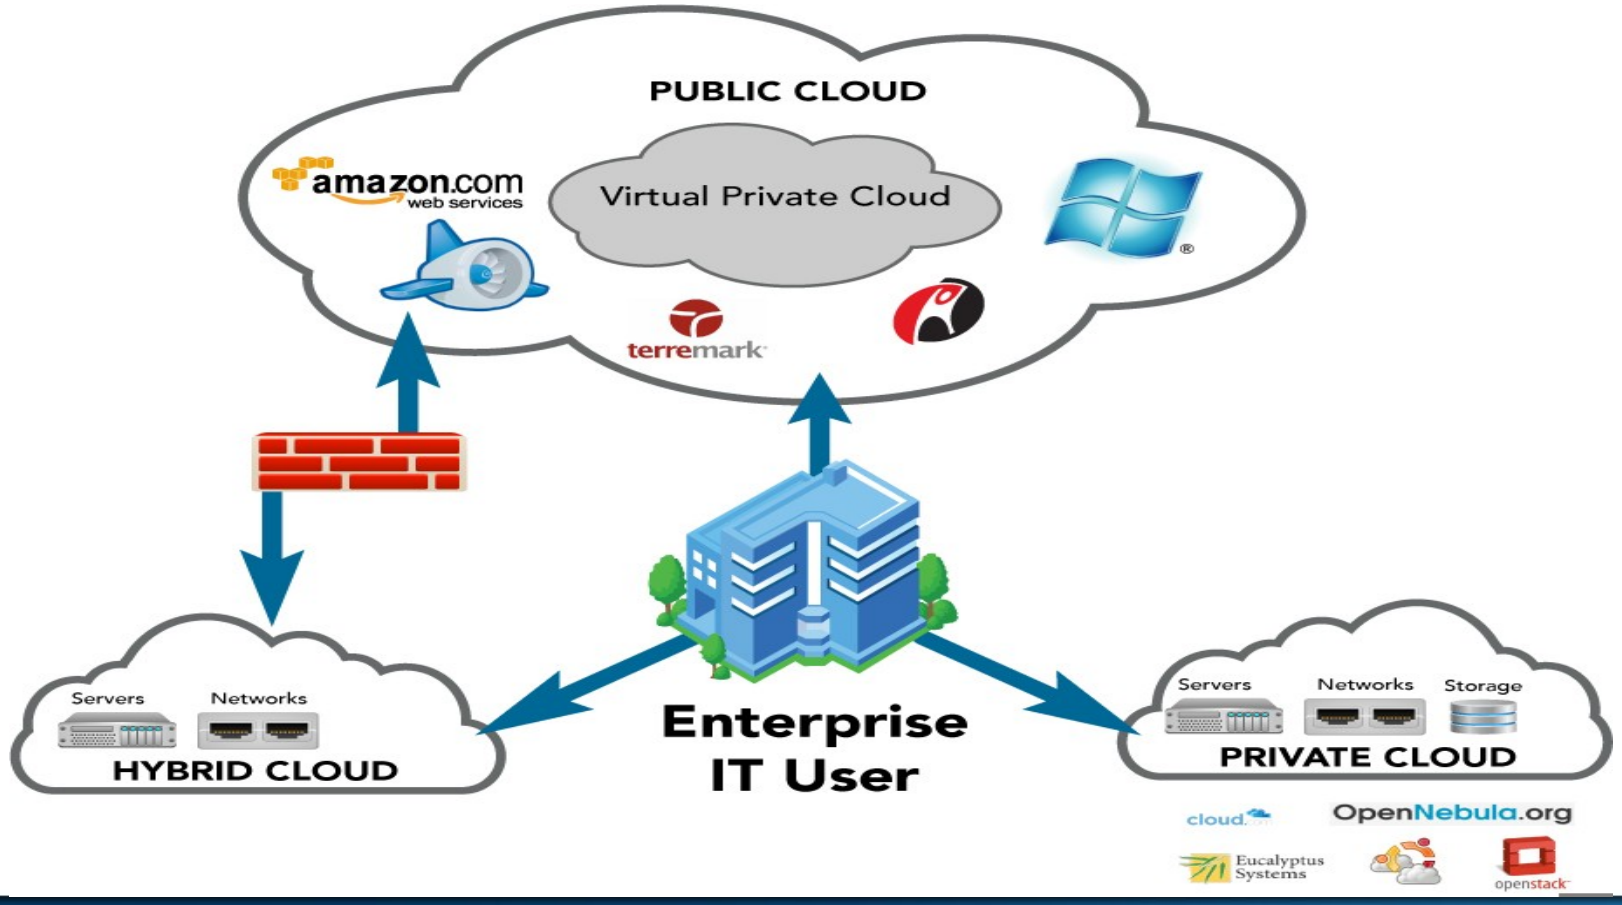

Building Compute Clouds with Open Source Software

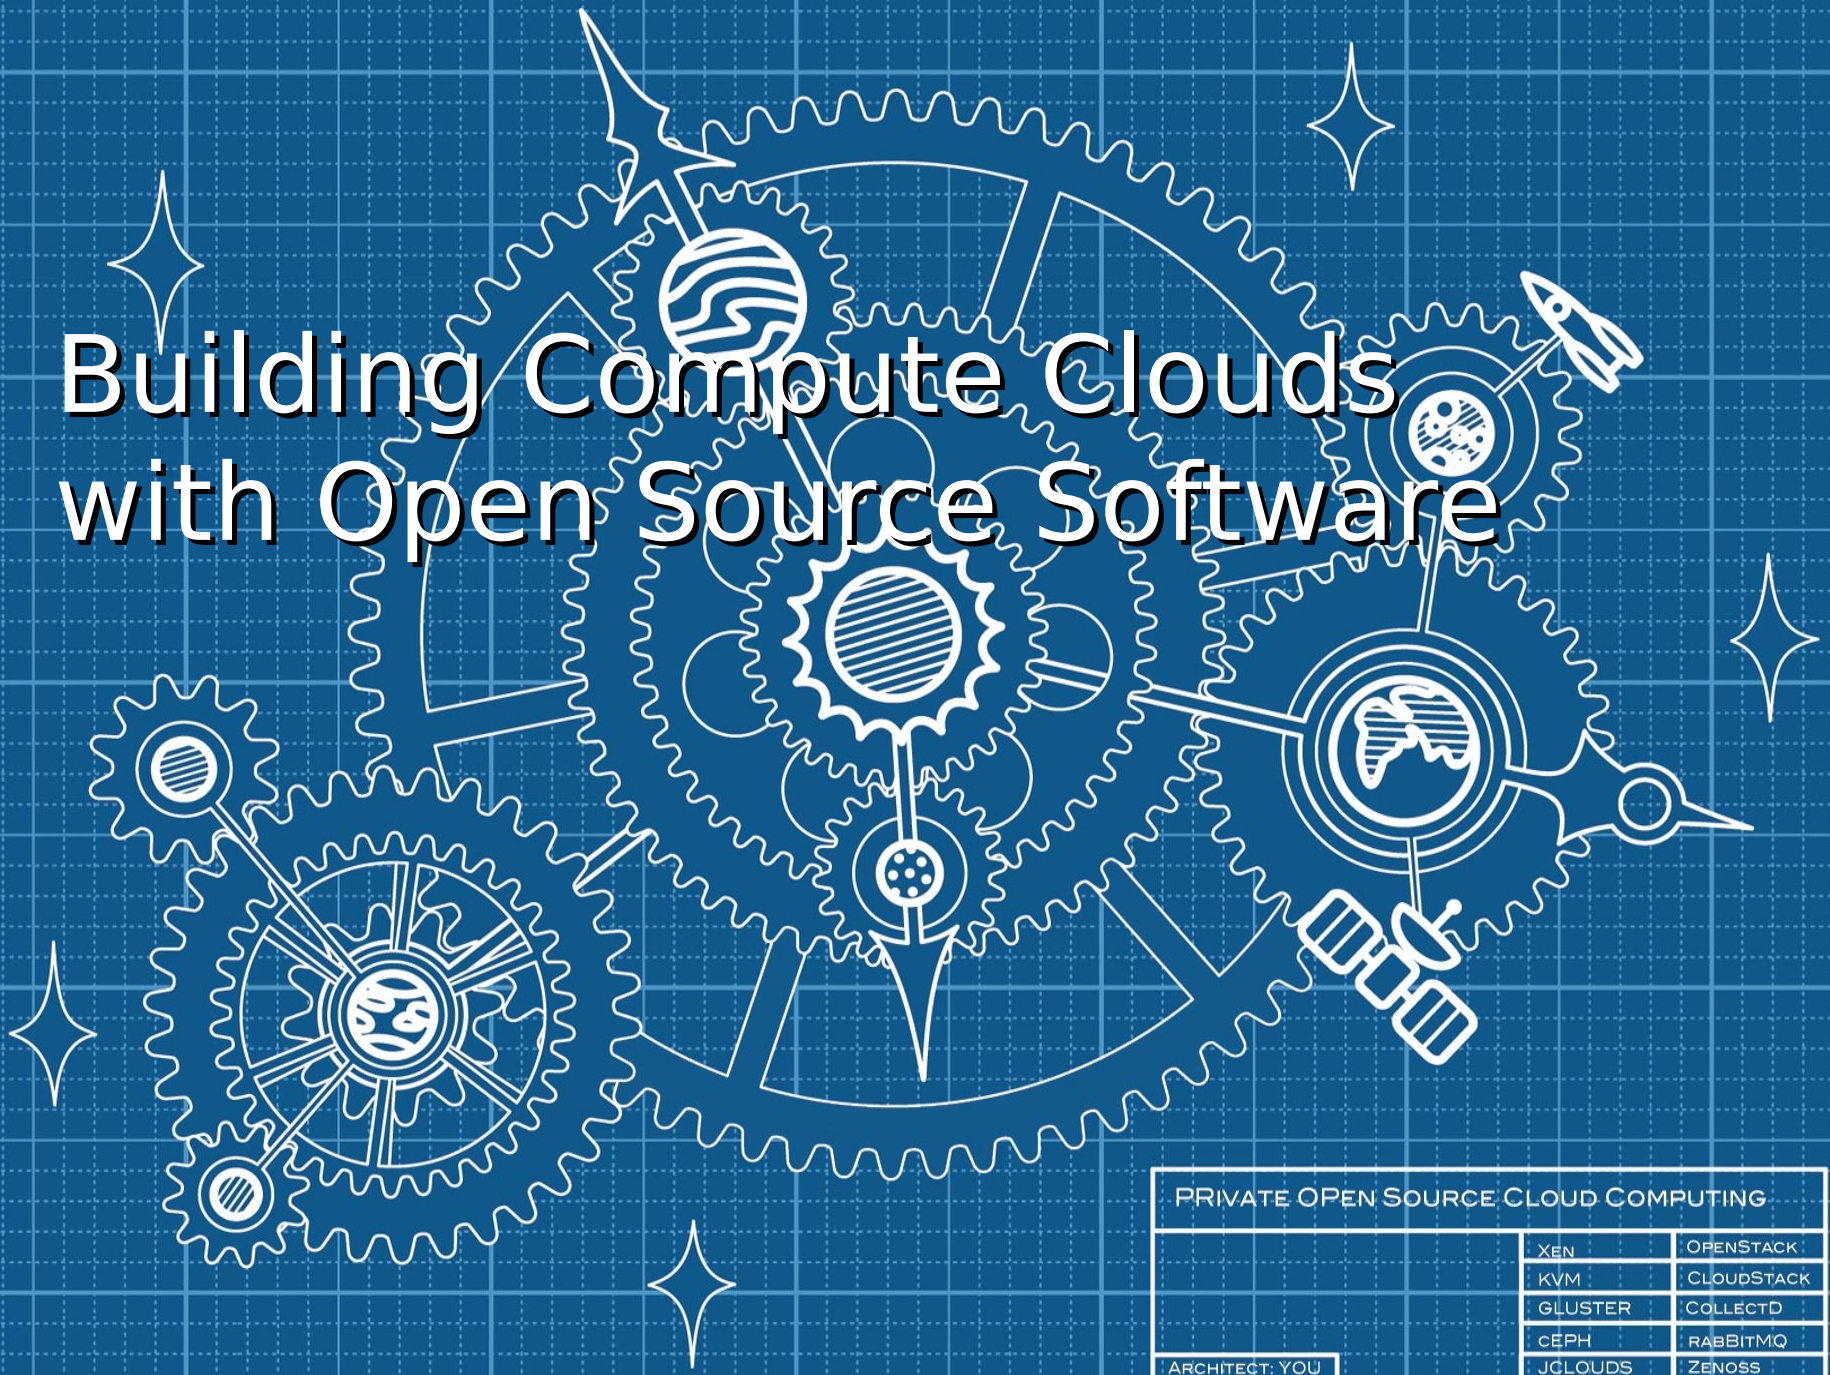

Cloud Still Requires Architectural Design

- Cloud Computing isn't a "magical solution"
- Need to design your architecture with *the end in mind*
- As you build it make your infrastructure easily *replicable*

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 Zones
- LXC – Userspace chrooted installs

Proprietary

- VMware
- Citrix Xenserver
- Microsoft Hyper-V
- Oracle VM

OpenVZ

Open Source Compute Clouds

	Year Started	License	Hypervisors Supported
cloud.com	2010 (Development Since 2008)	GPLv3+	Xenserver, XCP, VMware, KVM
Eucalyptus Systems	2008	GPLv3	Xen, KVM, VMware
openstack	2010 (Developed by NASA by Anso Labs previously)	Apache	Xen, KVM, Hyper-V
Ubuntu Enterprise Cloud	2009 (Karmic Koala)	GPLv3	Xen,KVM
abiquo	2009 (Development 2006)	LGPLv3	VMware ESX and ESXi, Microsoft Hyper-V, Xen, KVM and Virtual Box

Open Source Platform-as-a-Service

	Year Started	Sponsors	Platforms Supported
CLOUD FOUNDRY	2011	VMware	Spring, Rails, Sinatra, Node.js
OPENSHIFT PaaS by Red Hat Cloud	2011	Red Hat	Express – Ruby, PHP Python Flex – JBoss, Java EE6
WSO2 Stratos	2009	WS02	Tomcat, JBoss, Java EE6,

Open Source Cloud Computing Storage

- GlusterFS – Scale Out NAS system aggregating storage over Ethernet or Infiniband
- CloudFS – GlusterFS, with multi-tenant, encryption, additional management support
- CEPH – Highly capable distributed file storage system
- 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

- jclouds
- libcloud
- deltacloud
- fog

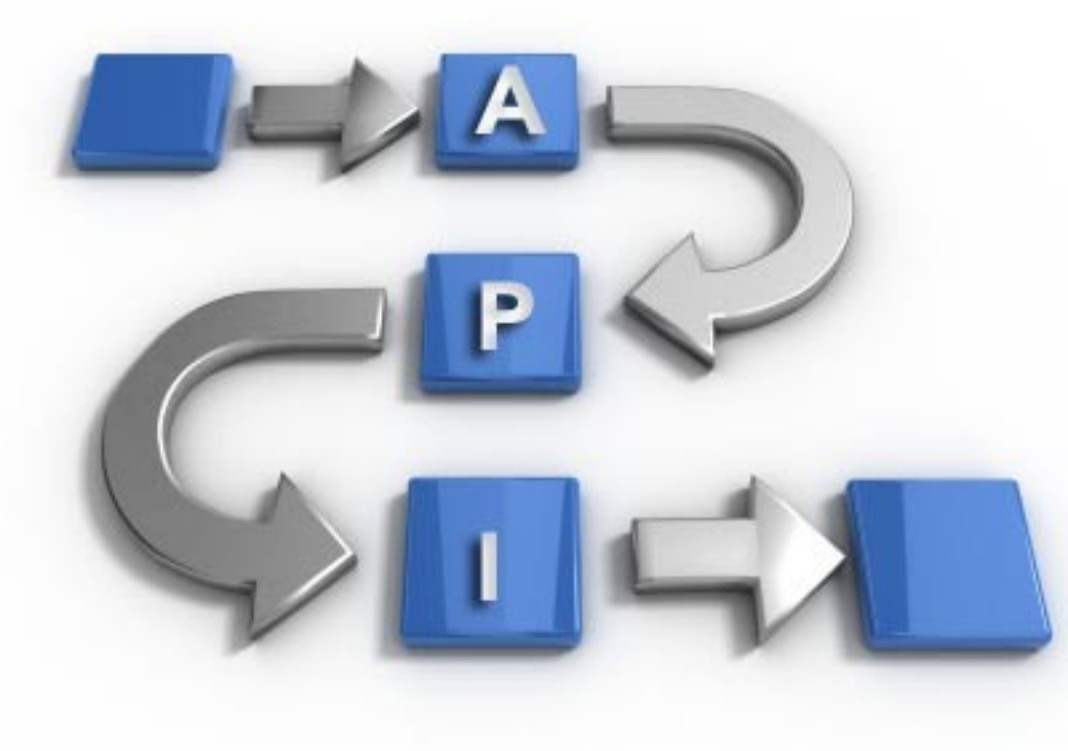

Managing Clouds with Open Source Tools

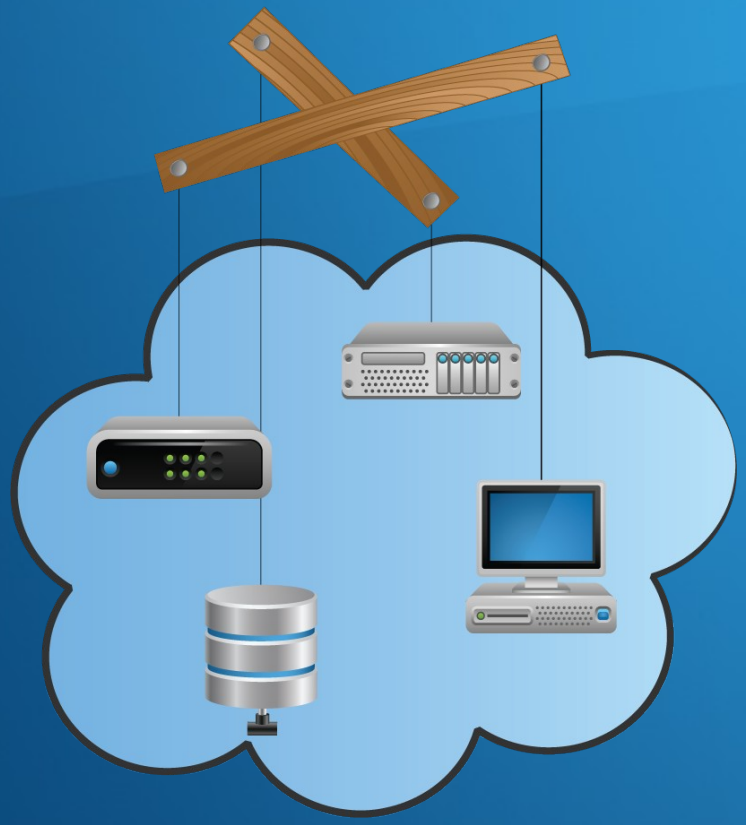

cloud.com
Open Source Cloud Computing Software

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*

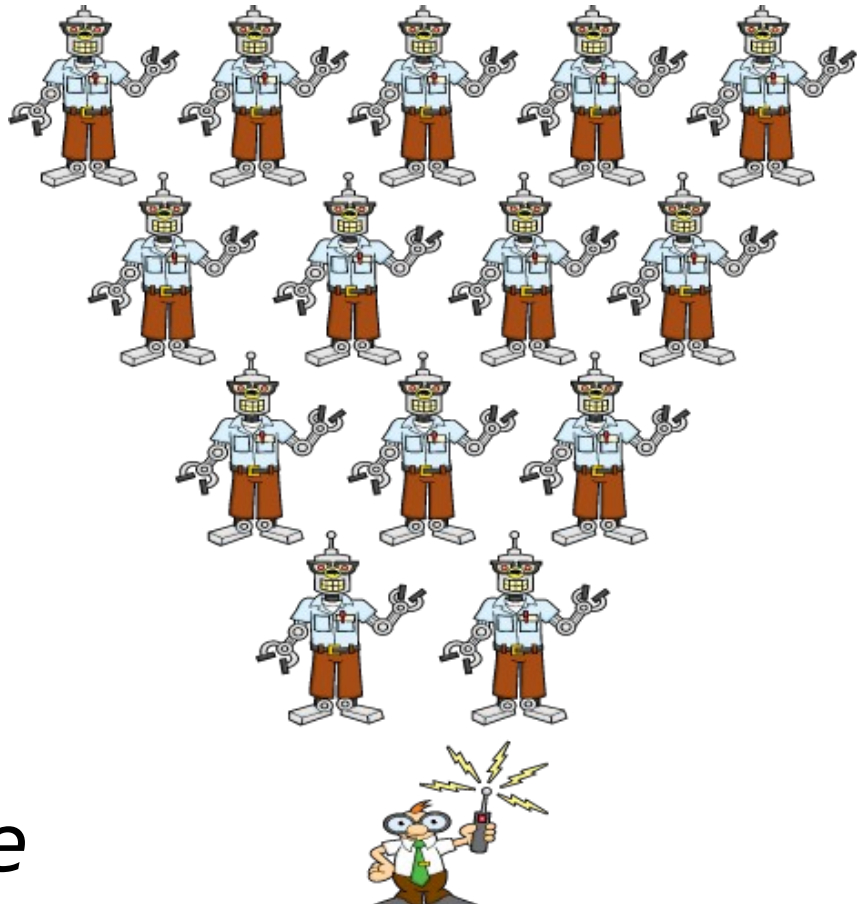

Why Open Source Tools?

- Aggressively Developed, Keep Pace with State of the Art
- User-Developed and Instrumented
- Easy to Assemble into Automated Toolchains

4 Types of Management Tools

Provisioning
Installation of operating systems and other software

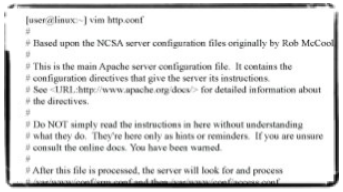

Configuration Management
Sets the parameters for servers, can specify installation parameters

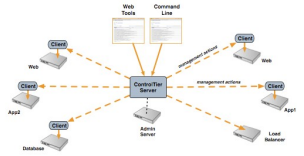

Orchestration/Automation
Automate tasks across systems

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

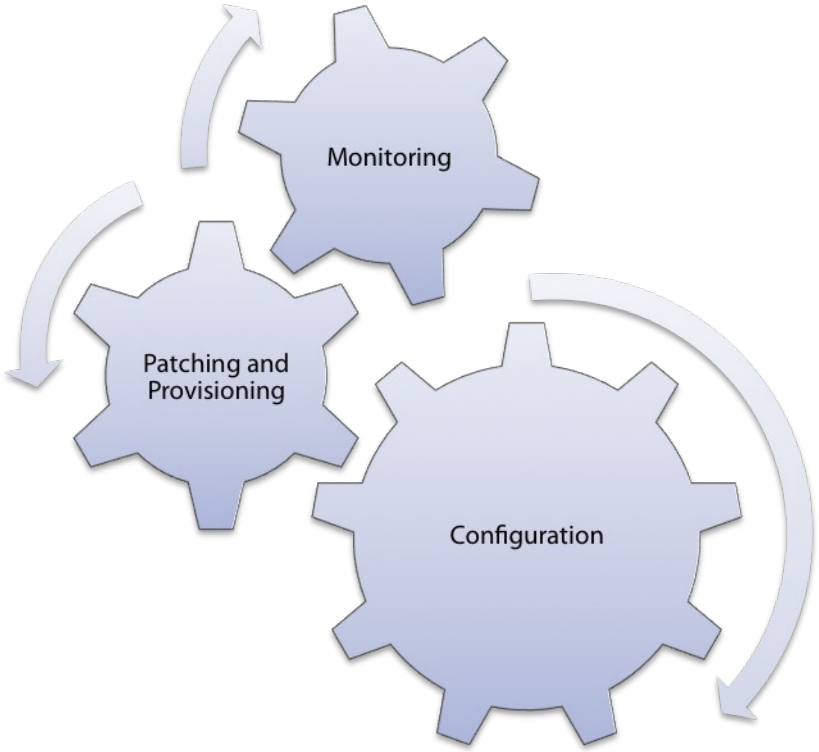

Open Source Provisioning Tools

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

Open Source Configuration Management Tools

	Year Started	Language	License	Client/Server
Bcfg2	2003	Python	BSD	Yes
Cfengine	1993	C	Apache	Yes
Chef	2009	Ruby	Apache	Chef Solo – No Chef Server - Yes
Puppet	2004	Ruby	GPL	Yes

Open Source Monitoring Tools

	Year Started	License	Language	Type of Monitoring	Collection Methods
Cacti	2001	GPL	PHP	Performance	SNMP, syslog
Nagios	1999	GPL	C/PHP	Availability	SNMP,TCP, ICMP, IPMI, syslog
OpenNMS	2000	GPL	Java	Availability/Performance	SNMP,
Zabbix	2001	GPL	C/PHP	Availability/ Performance and more	SNMP, TCP/ICMP, IPMI, Synthetic Transactions
Zenoss	2005	GPL	Python	Availability, Performance, Event Management	SNMP, ICMP, SSH, syslog, WMI

Open Source Automation/Orchestration Tools

	Year Started	Language	License	Client/Server	Support Organization
AutomateIT	2009	Ruby	GPL	No	None
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

Automated Toolchain

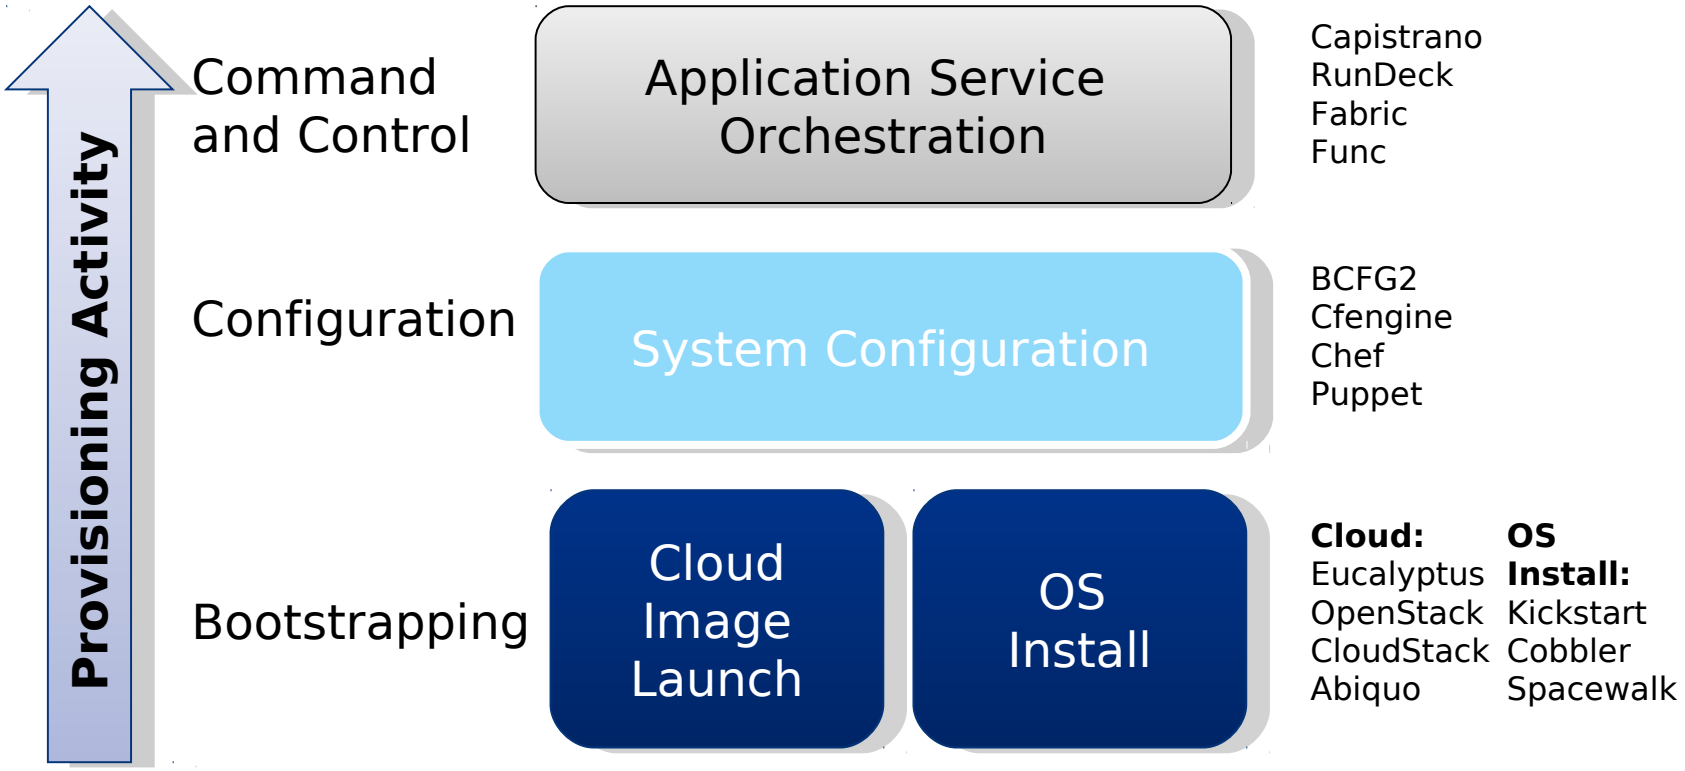

Questions?

Contact Me

Professional:david.nalley@cloud.com
Personal: david@gnsa.us

IRC: ke4qqq on irc.freenode.net
normally hanging out on #cloudstack, #fedora-* and 50 other channels.

Professional: http://cloudstack.org
Personal: http://ke4qqq.wordpress.com

Identi.ca/twitter: @ke4qqq

Crash Course in Open Source Cloud Computing