Hitchhiker's Guide to Open Source Cloud Computing

By Mark R. Hinkle
Senior Director, Cloud Computing Community
Citrix Inc.
Twitter: @mrhinkle
Blog: www.socializedsoftware.com
Email: mrhinkle@socializedsoftware.com

%whoami

- Dedicated to the success of the Apache CloudStack & Xen Open Source Cloud Communities at Citrix
- Conduct Build A Cloud learning activities all over the world
- Joined Citrix via Cloud.com acquisition July 2011
- 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 Open Source Subnet

Why Open Source and the Cloud?

- User-Driven Context from Solving Real Problems
- Lower Barrier to Participation
- Larger user base, users helping users
- Aggressive release cycles stay current with the state-of-the-art
- Open Source innovating faster than commercial
- Open data, Open standards, Open APIs

Quick Cloud Computing Overview or the Obligatory “What is the Cloud Explanation”

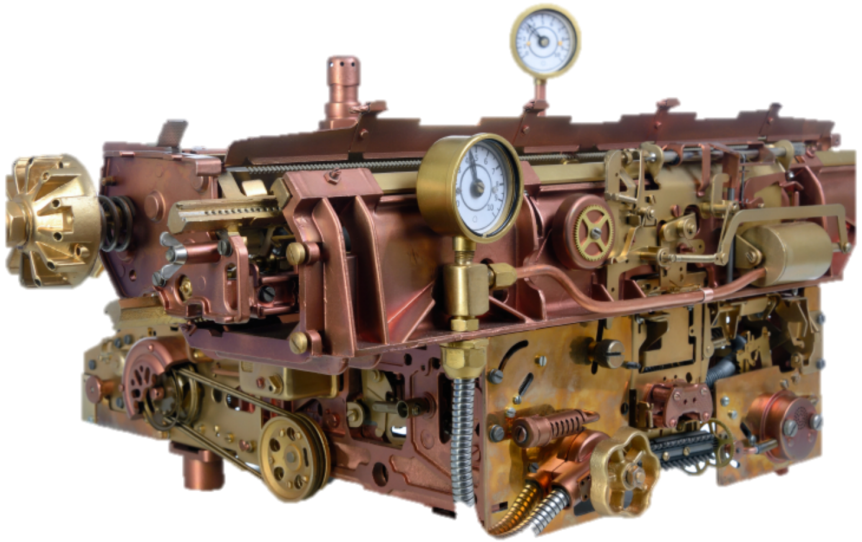

Infinite Probability Drive

Five Characteristics of Cloud

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

Cloud Computing Service Models

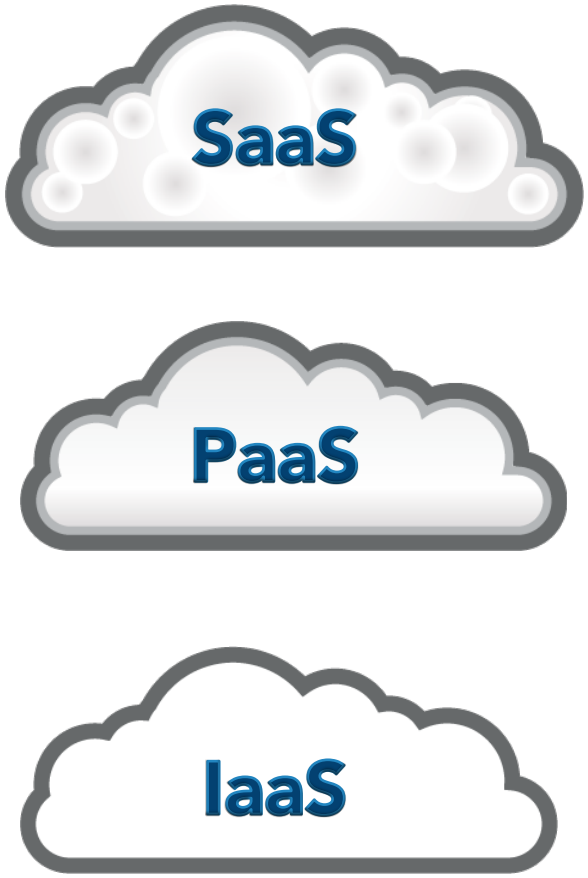

USER CLOUD a.k.a. SOFTWARE AS A SERVICE

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: 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: Public, Private & Hybrid

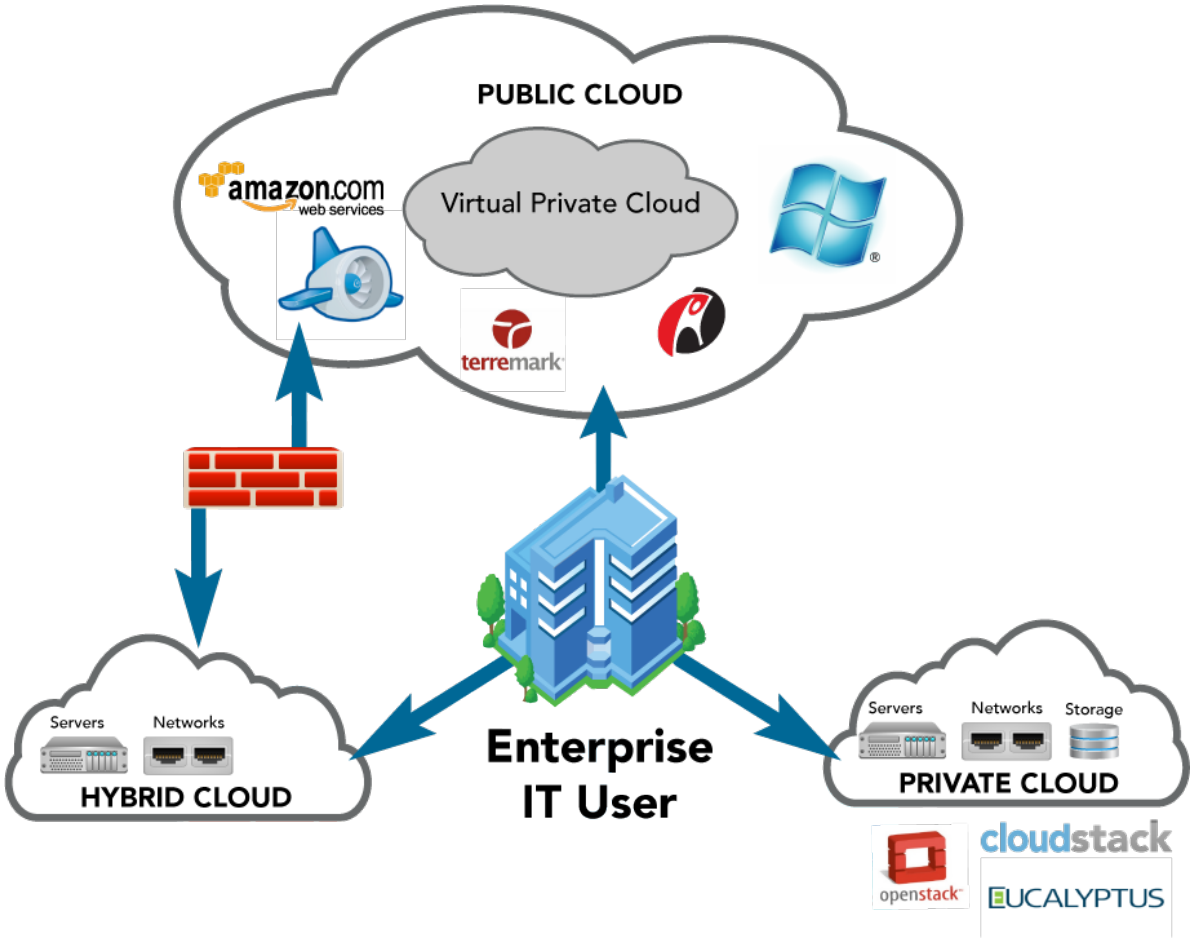

Building Open Source Clouds

Cloud Architecture

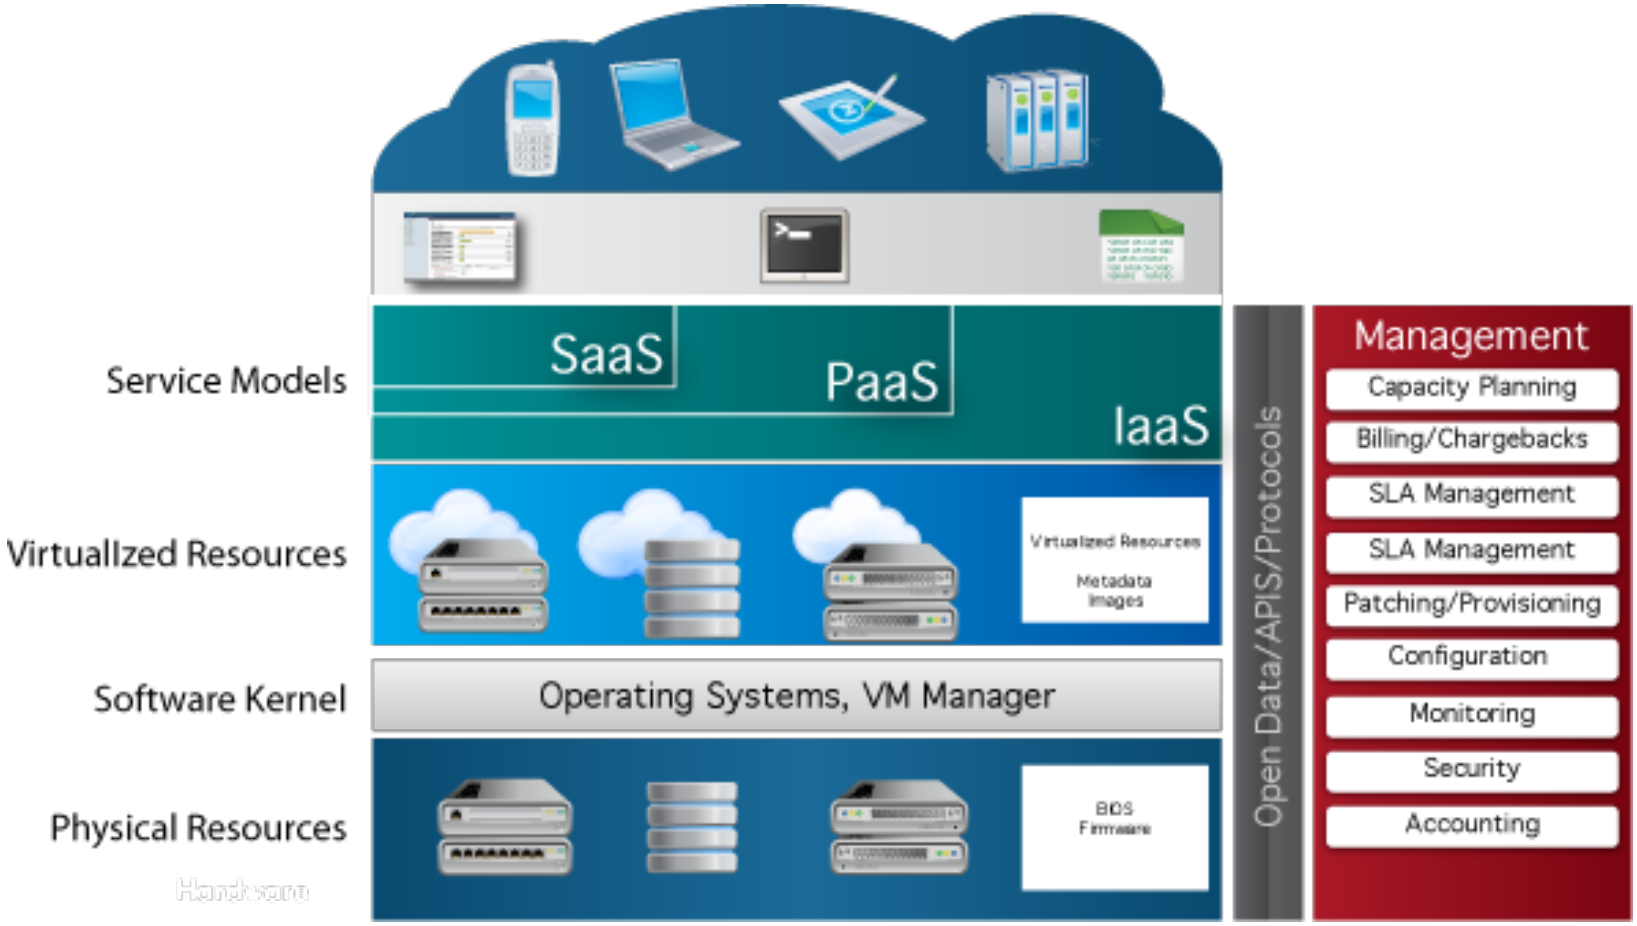

Hypervisors

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 (based
- Microsoft Hyper-V
- OracleVM (Based on OS Xen)

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 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
Oz	Command-line tool that has the ability to create images for common Linux distributions to run on KVM
SUSE Studio	SUSE Studio supports building and deploying directly to cloud services such as Amazon EC2.

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)

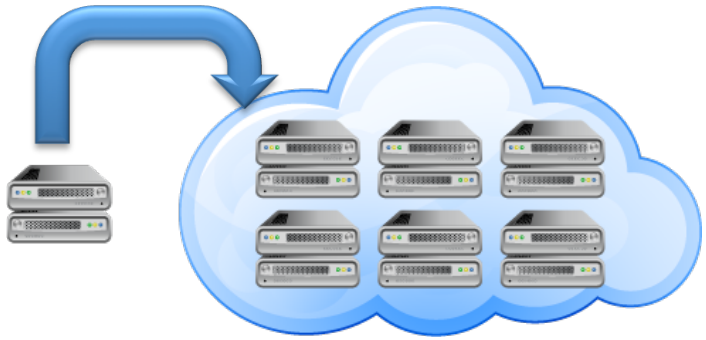

Compute Clouds (IaaS)

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

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

OpenStack – Ecosystem of Projects

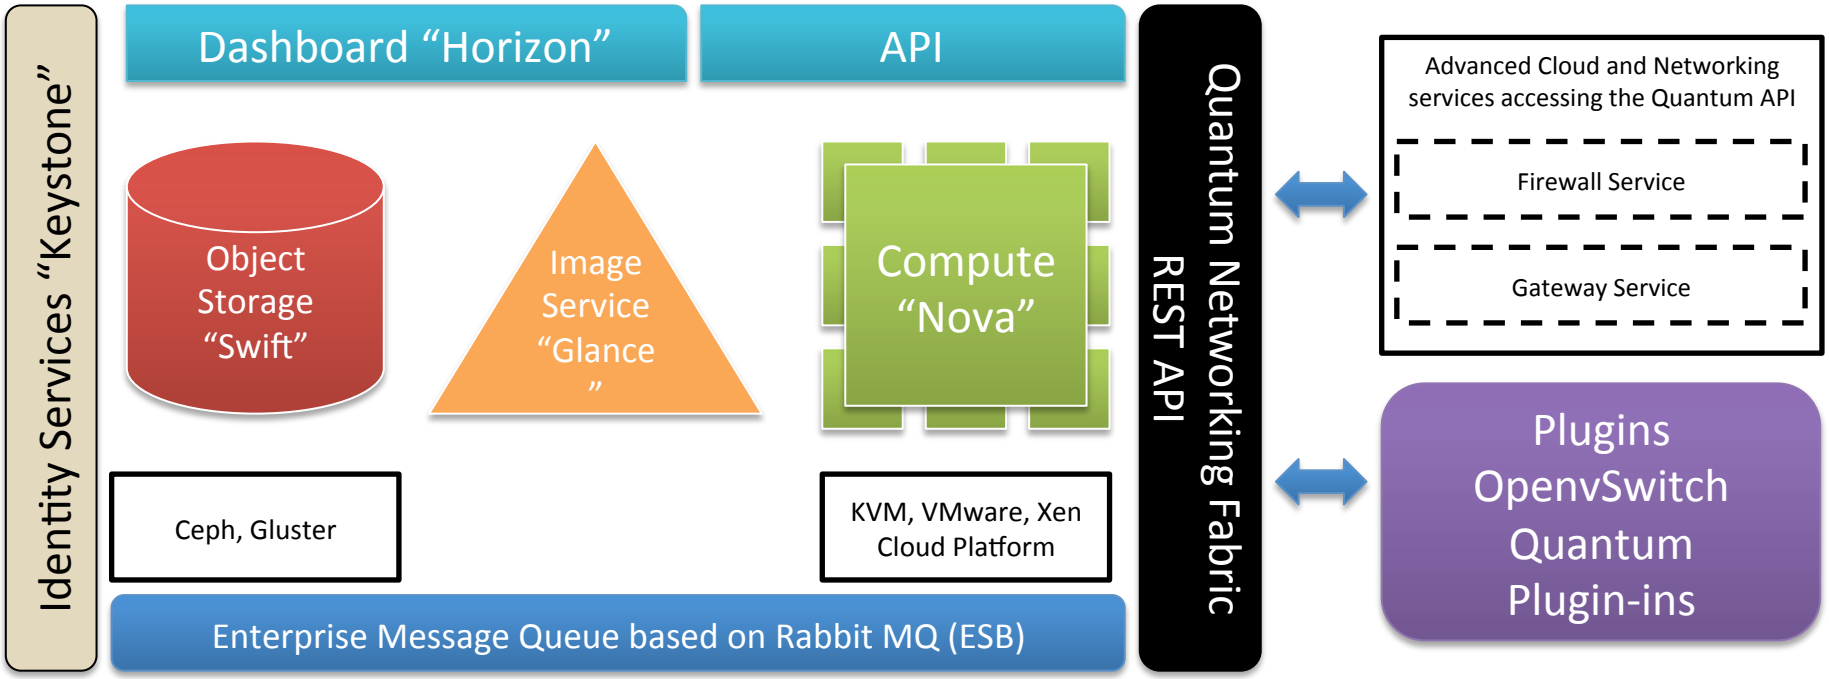

20 Collective projects hosted at:

https://launchpad.net/openstack

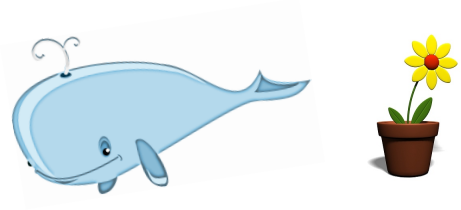

Cloud APIs

- jclouds
- libcloud
- deltacloud
- fog

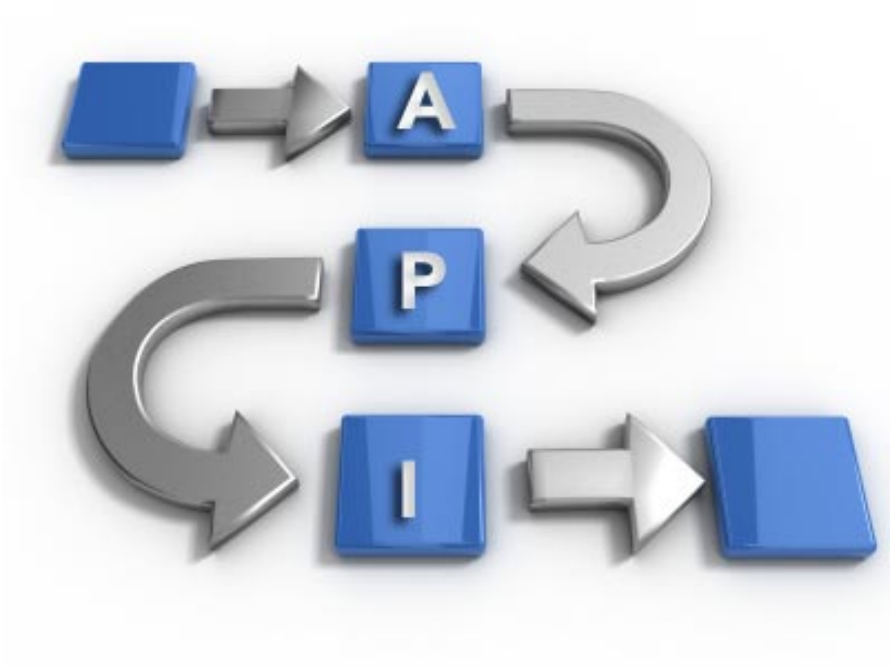

Cloud Computing Storage

Project	Description
GlusterFS	Scale Out NAS system aggregating storage over Ethernet or Infiniband
Ceph	Distributed file storage system developed by DreamHost
OpenStack Storage	Long-term object storage system
Sheepdog	Distributed storage for KVM hypervisors

Platform-as-a-Service (PaaS)

Project	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)
Cloudify	2012	Gigaspaces	
OpenShift **	2011	Red Hat	Java, Ruby, PHP, Perl and Python
Stackato*	2012	ActiveState	Java, Python, PHP, Ruby, Perl, Node.js, others
WSO2 Stratus	2010	WSO2	Jboss, Java EE6

Software Defined Networking (SDN)

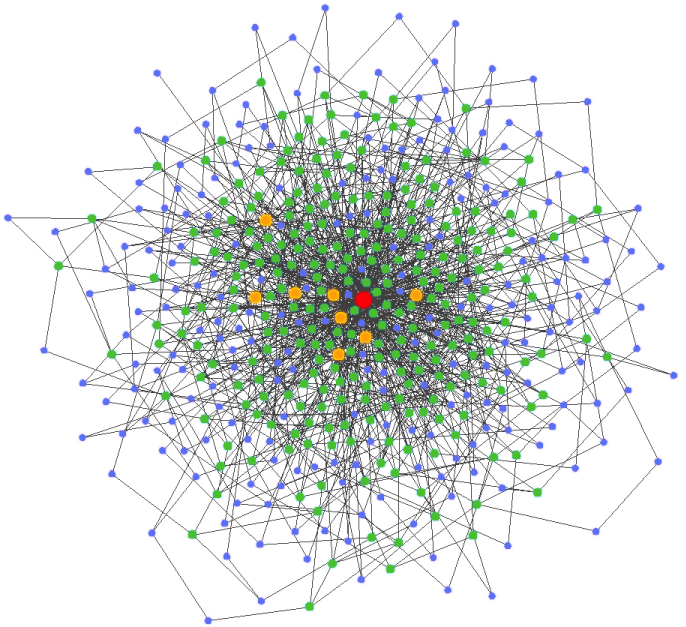

Overview of Software Defined Networking

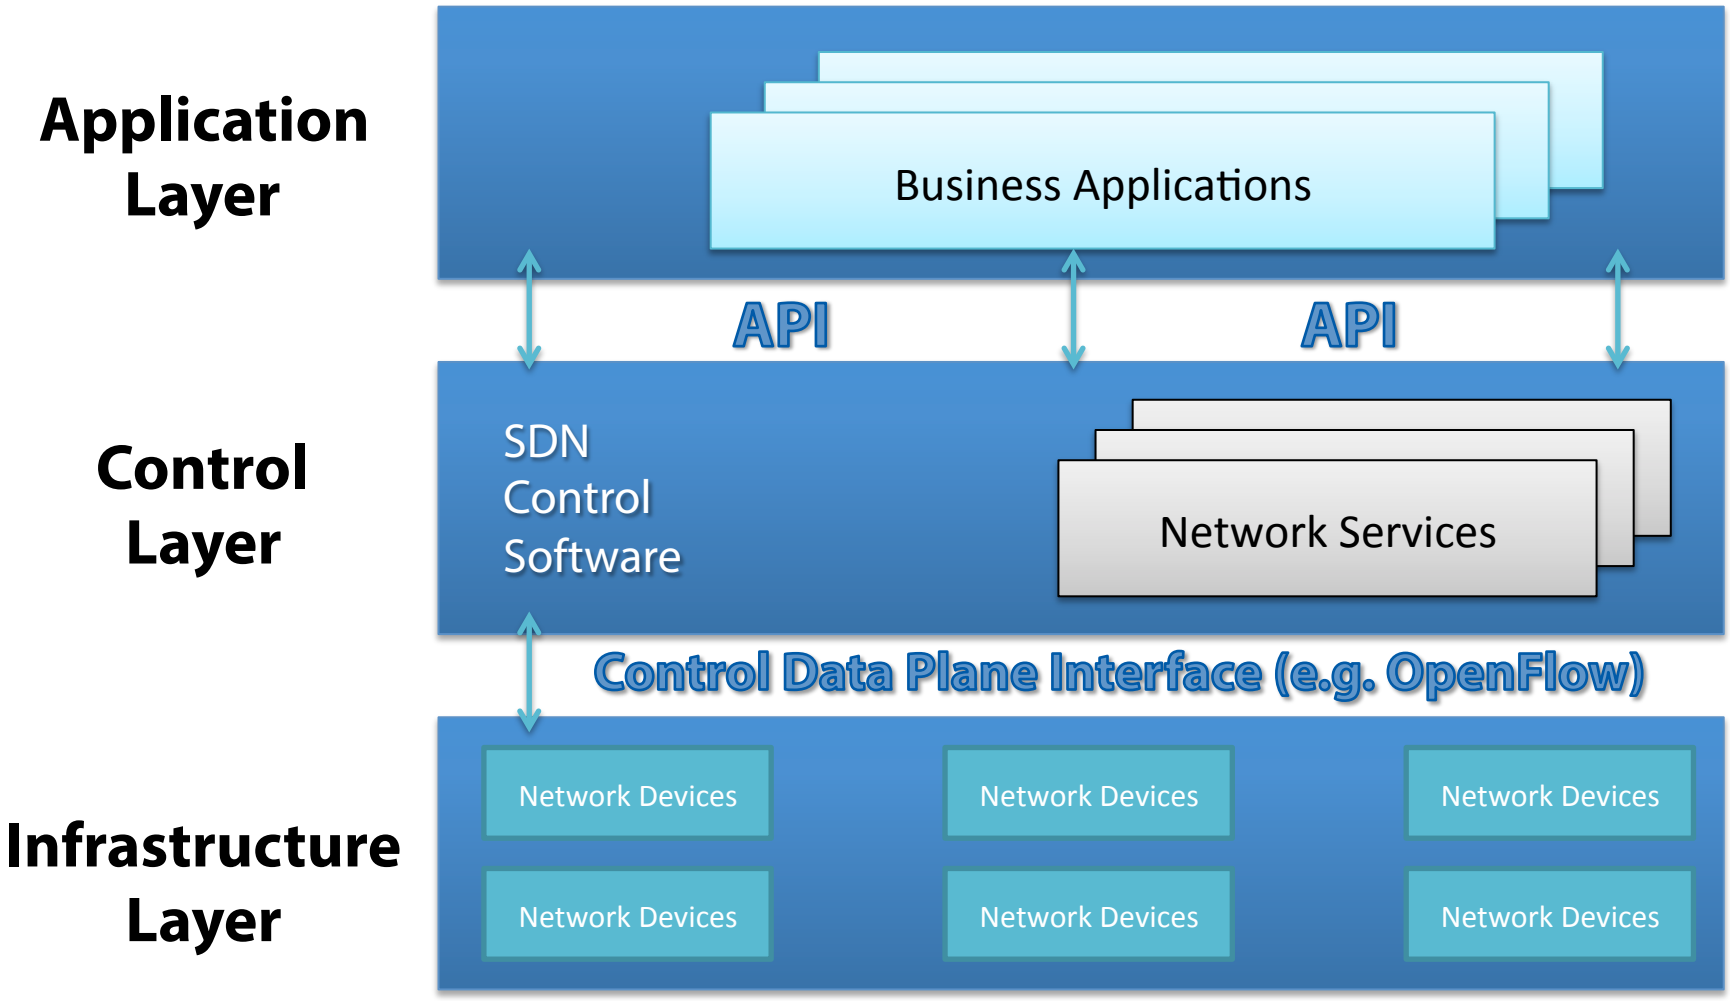

Why SDN?

Cloud Promise	Cloud Reality
Centralized Configuration and Automation	Without true virtualization, network devices must still be manually configured.
Instant Self-Service Provisioning	In a physical network, it could take a long time for network engineer to provision new services.
Elasticity and Scalability	By horizontally scaling up the physical network, elasticity is lost.
Designed for Failure	Failover can be automated and physical network limitations can be alleviated.

Source: Midokura

Open Flow

OpenFlow enables networks to evolve, by giving a remote controller the power to modify the behavior of network devices, through a well-defined "forwarding instruction set". The growing OpenFlow ecosystem now includes routers, switches, virtual switches, and access points from a range of vendors.

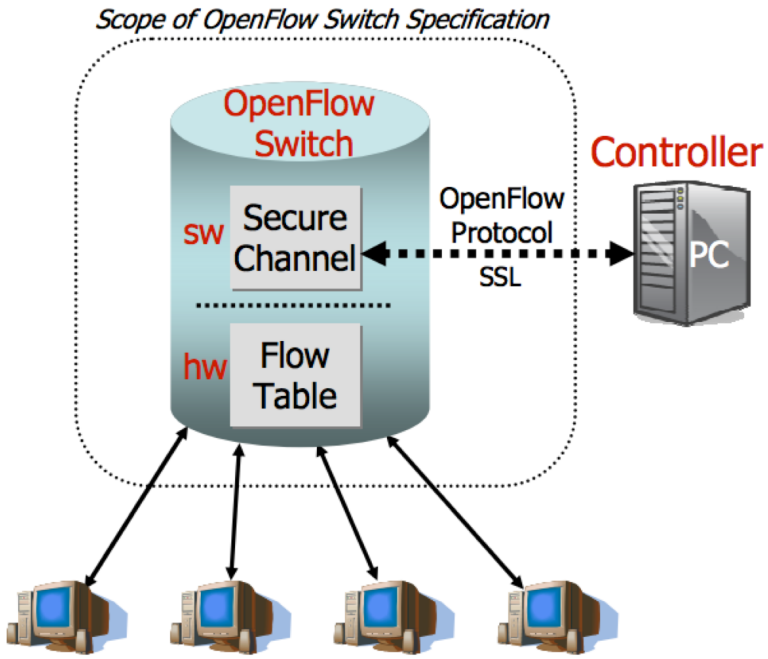

Image from http://www.openflow.org/documents/openflow-wp-latest.pdf

Software Defined Networking

Project	Description
Floodlight	The Floodlight controller is an enterprise-class, Apache-licensed, Java-based OpenFlow Controller.
Indigo	Indigo is an open source project to support OpenFlow on a range of physical switches. By leveraging hardware features of Ethernet switch ASICs, Indigo supports high rates for high port counts, up to 48 10-gigabit ports. Multiple gigabit platforms with 10-gigabit uplinks are also supported.
OpenStack Networking "Quantum"	Pluggable, scalable, API-driven network and IP management
Open vSwitch	Open vSwitch is a open source (ASL 2.0), multilayer virtual switch designed to enable massive network automation through programmatic extension, while still supporting standard management interfaces and protocols (e.g. NetFlow, sFlow, SPAN, RSPAN, CLI, LACP, 802.1ag).

Big Data

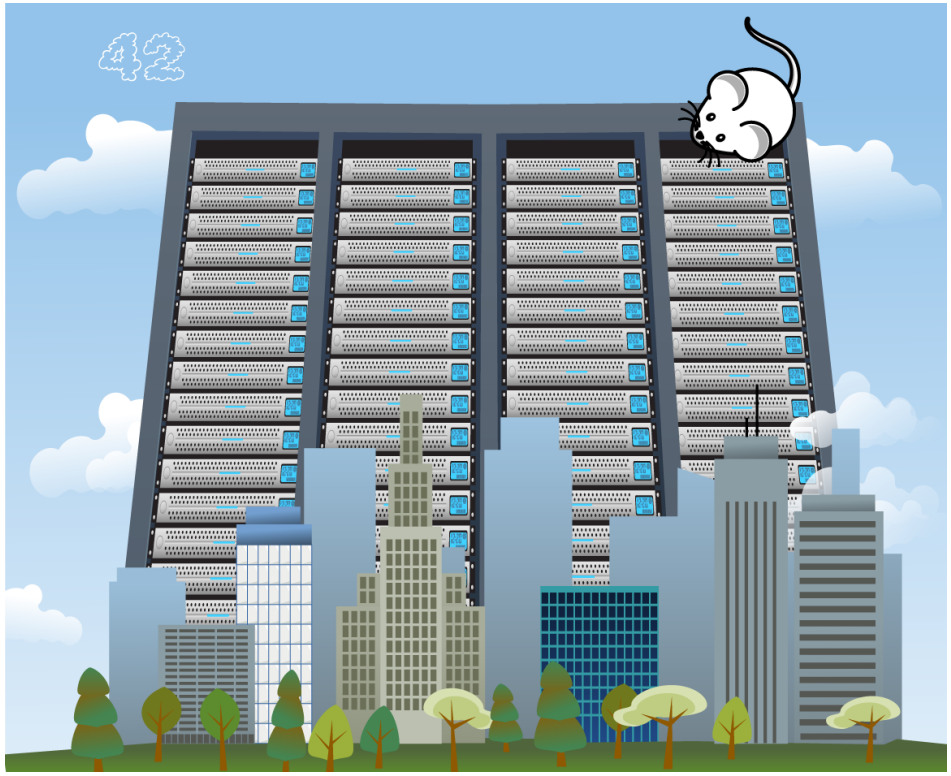

Deep Thought

1 Billion Facebook Users - October 2012

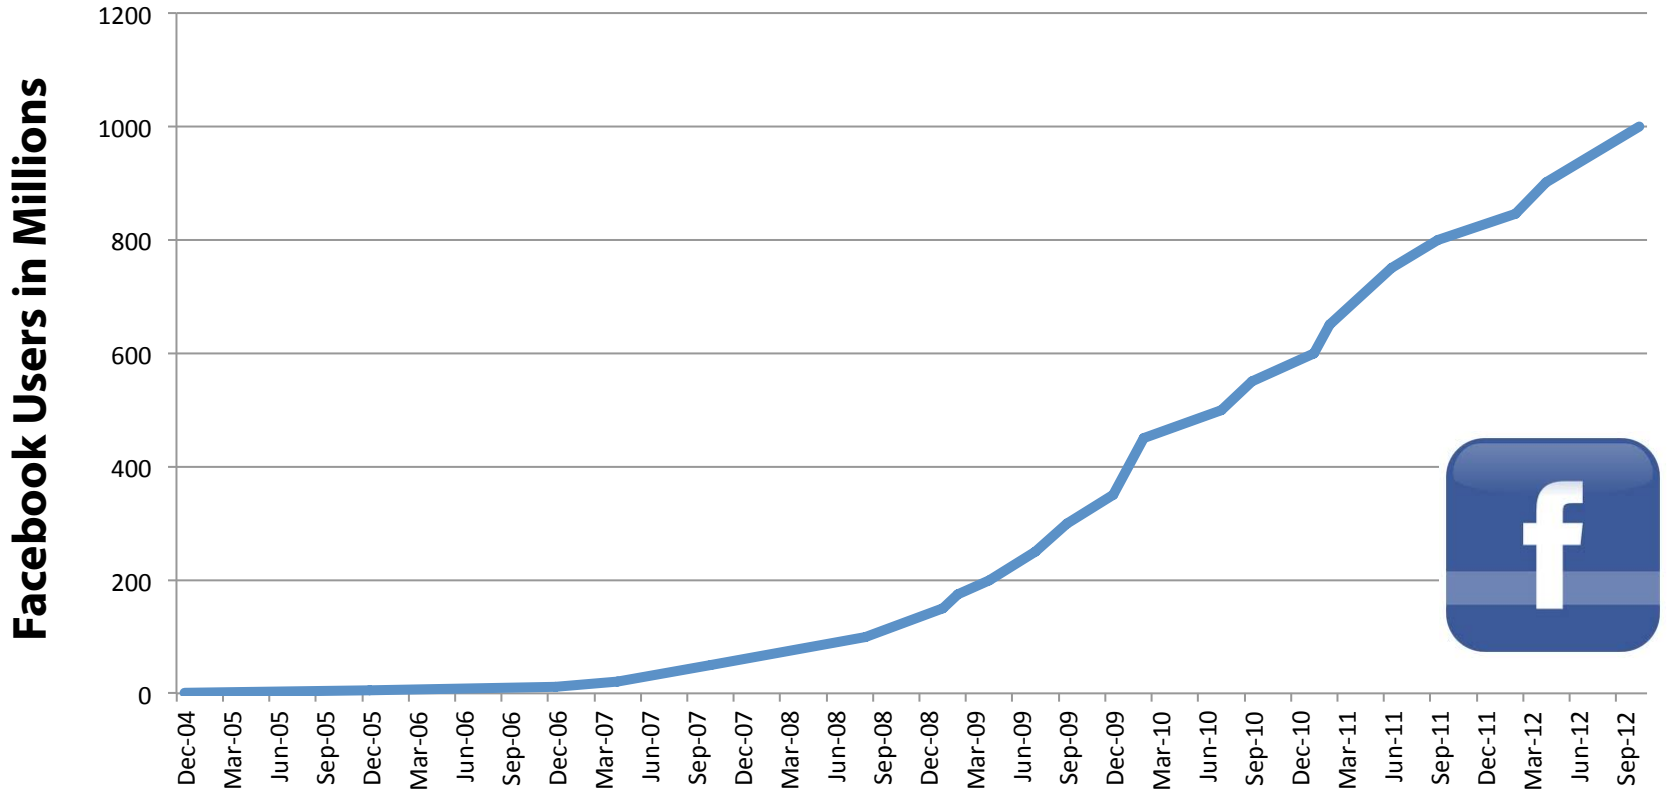

Source: Benphoster.com

Twitter at 400M Tweets Per Day – June 2012

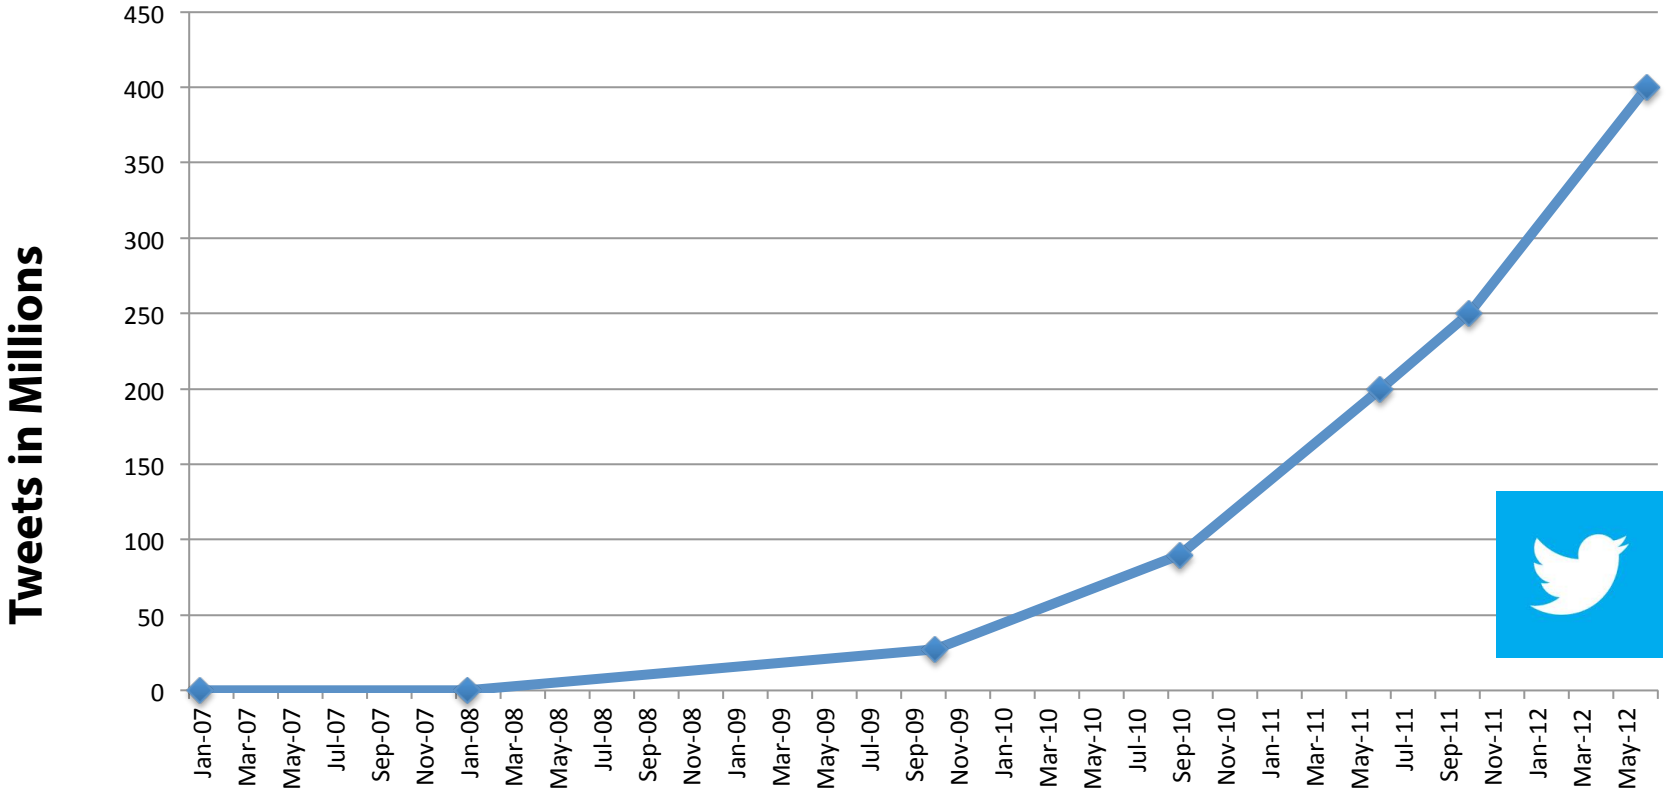

Source :TheBigDataGroup.com

Big Data and Storage Infrastructure

Data is growing faster than storage capacity and computing power. Legacy systems hold organizations back; storage software must include multi-petabyte capacity, support potentially billions of objects, and provide application performance awareness and agile provisioning.

-Gartner, Big Data Challenges for the IT Infrastructure Team

Big Data Landscape

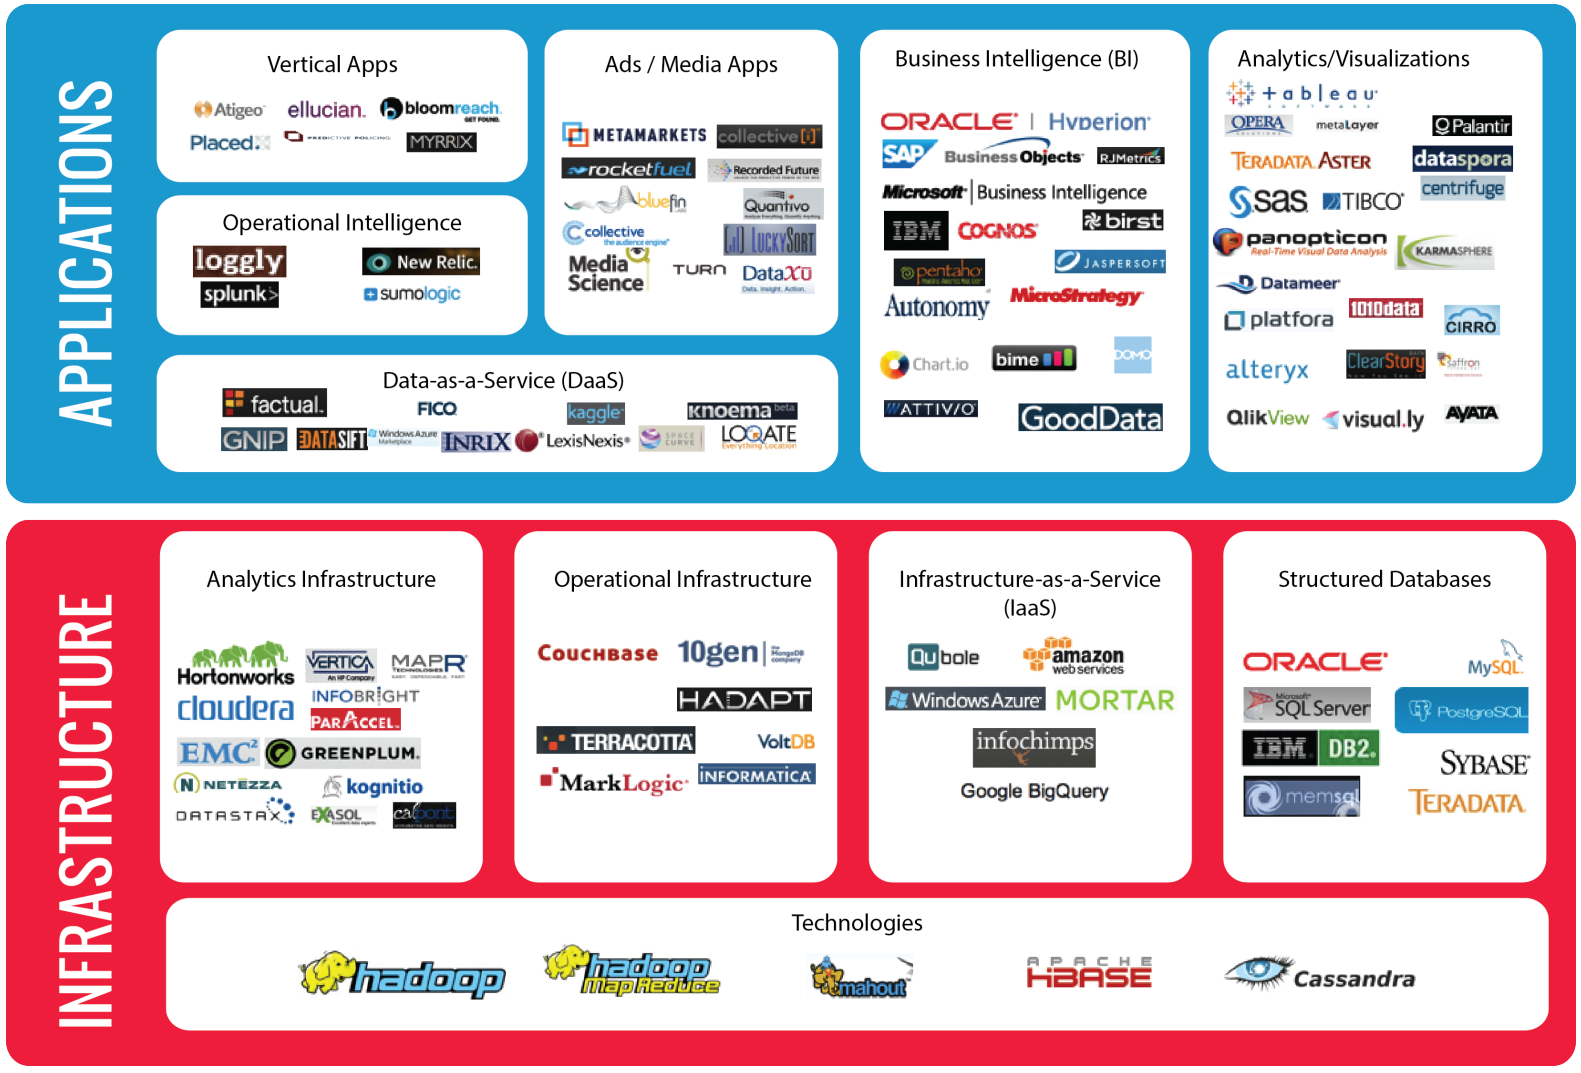

Source: BigDataGroup.com

Open Source NoSQL Databases

Name	Type	Description
Apache Cassandra	Wide Column Store/ Families	API: many » Query Method: MapReduce, Replicaton: , Written in: Java, Concurrency: eventually consistent , Misc: like "Big-Table on Amazon Dynamo alike", initiated by Facebook
CouchDB	Document Store	API: Memcached API+protocol (binary and ASCII) , most languages, Protocol: Memcached REST interface for cluster conf + management, Written in: C/C++ + Erlang (clustering), Replication: Peer to Peer, fully consistent, Misc: Transparent topology changes during operation, provides memcached-compatible caching buckets
HBase	Wide Column Store/ Families	API: Java / any writer, Protocol: any write call, Query Method: MapReduce Java / any exec, Replication: HDFS Replication, Written in: Java
Hypertable	Wide Column Store/ Families	PI: Thrift (Java, PHP, Perl, Python, Ruby, etc.), Protocol: Thrift, Query Method: HQL, native Thrift API, Replication: HDFS Replication, Concurrency: MVCC, Consistency Model: Fully consistent Misc: High performance C++ implementation of Google's Bigtable.
MongoDB	Document Store	API: BSON, Protocol: C, Query Method: dynamic object-based language & MapReduce, Replication: Master Slave & Auto-Sharding, Written in: C++,Concurrency
Redis	Key Value/ Tuple Store	API: Tons of languages, Written in: C, Concurrency: in memory and saves asynchronous disk after a defined time. Append only mode available. Different kinds of fsync policies. Replication: Master / Slave, Misc: also lists, sets, sorted sets, hashes, queues.
Riak	Key Value / Tuple Store	API: JSON, Protocol: REST, Query Method: MapReduce term matching , Scaling: Multiple Masters; Written in: Erlang, Concurrency: eventually consistent (stronger then MVCC via Vector Clocks)

MapReduce

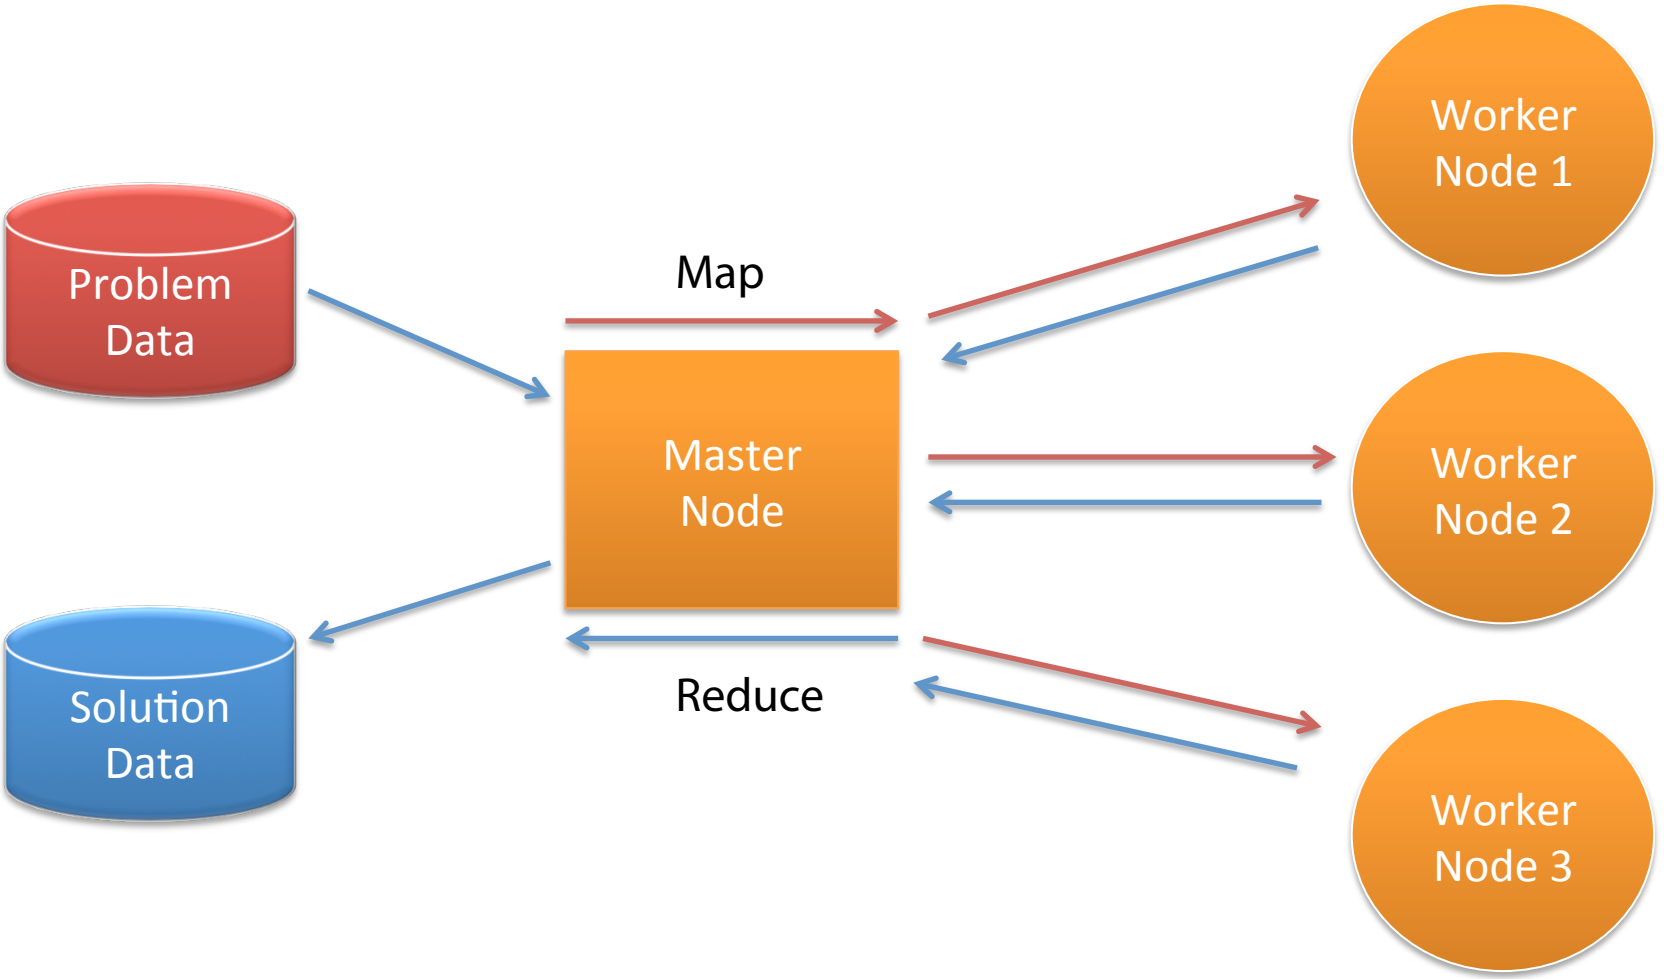
Problem Data
Solution Data
Master Node
Map
Reduce
Worker Node 1
Worker Node 2
Worker Node 3

Apache Hadoop

Overview

- Handles large amounts of data
- Stores data in native format
- Delivers linear scalability at low cost
- Resilient in case of infrastructure failures
- Transparent application scalability

Facts

- Apache top-level open source project
- One framework for storage and compute
  - HDFS – Scalable storage in Hadoop Distributed File System (HDFS)
  - Compute via the MapReduce distributed processing platform
- Domain Specific Language (DSL) - Java

Hadoop Architecture

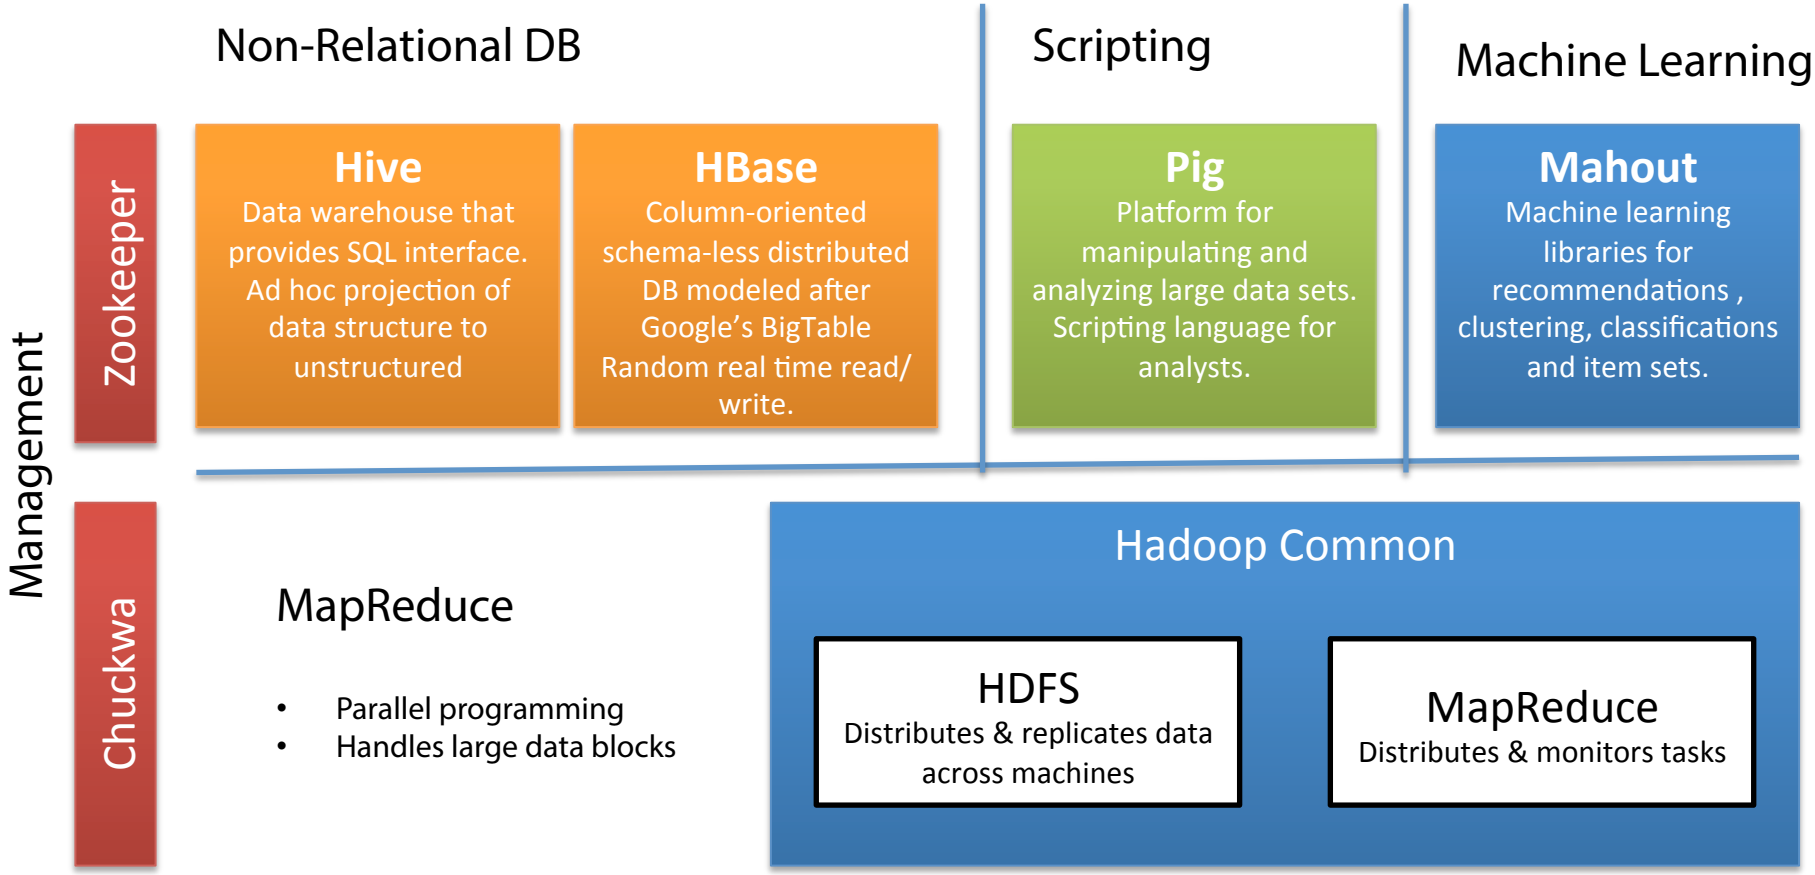

Big Data Summary

- Quantity of Machine Created Data Increasing Drastically (examples: networked sensor data from mobile phones and GPS devices)
- Data manipulation moving from batched to real-time
- Cloud services giving everyone Big Data tools
- Consumer company speed and scale requirements driving efficiencies in Big Data storage and analytics
- New and broader number of data sources being meshed together
- Big Data Apps means using Big Data is faster and easier

Cloud Management Tools

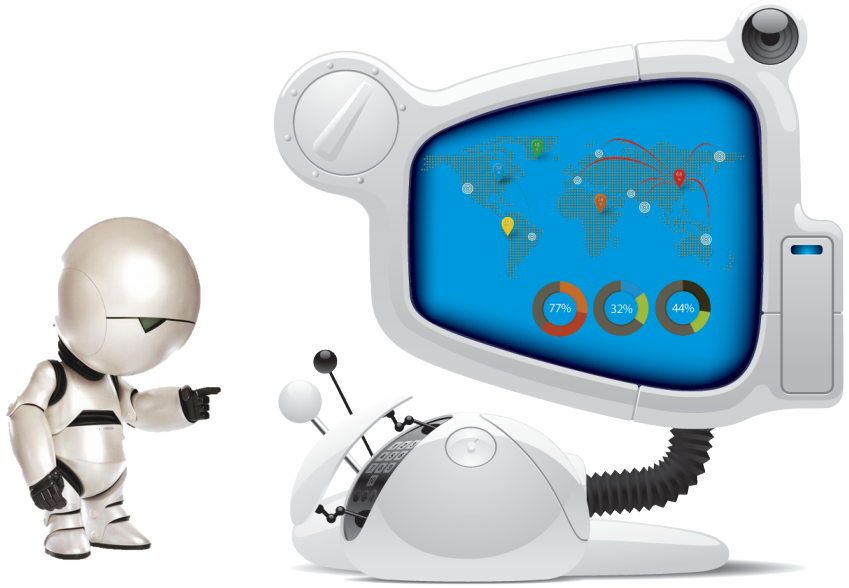

Automation in the Cloud

Meat Cloud

Cloud Operations

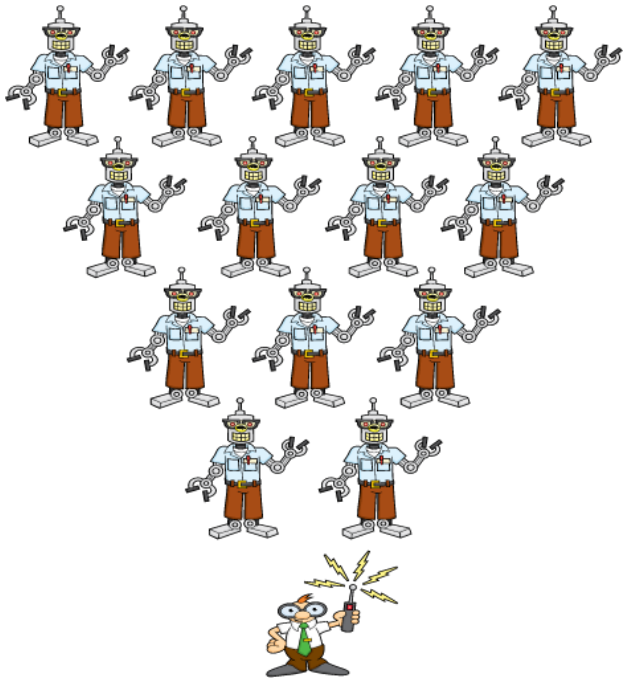

4 Types of Management Tools

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

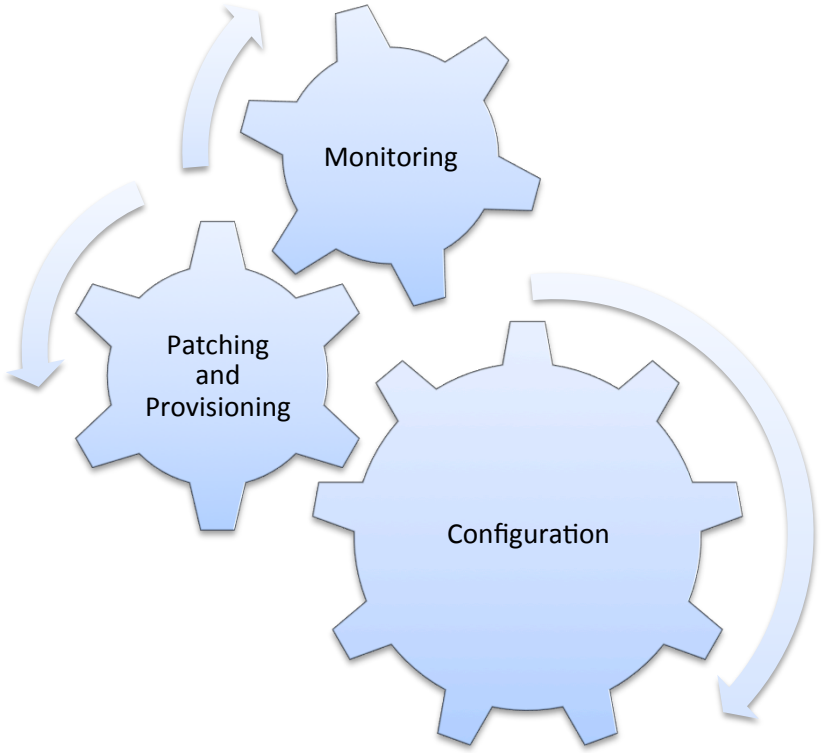

Provisioning

Project	Installation Targets
Axemblr Provisionr	Can provision 10s to 1000s of machines on various clouds.
Cobbler	Distributed virtual infrastructure using koan (kickstart of a network to PXE boot VMs) for Red Hat, OpenSUSE Fedora, Debian, Ubuntu VMs
JuJu	Public Clouds - Amazon Web Services HP Cloud, Private OpenStack clouds, Bare Metal via MAAS.
Salt Cloud	Tool to provision “salted” VMs that can then be updated by a central server via ZeroMQ
Crowbar	(Bare metal provisioning)

Configuration Management Tools

Project	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

Automation/Orchestration Tools

Project	Description
Capistrano	Utility and framework for executing commands in parallel on multiple remote machines, via SSH. It uses a simple DSL that allows you to define tasks, which may be applied to machines in certain roles
RunDeck	Rundeck is an open-source process automation and command orchestration tool with a web console.
Func	Func provides a two-way authenticated system for generically executing tasks, integrations with puppet and cobbler.
MCollective	The Marionette Collective AKA MCollective is a framework to build server orchestration or parallel job execution systems.
Salt	Execute arbitrary shell commands or choose from dozens of pre-built modules of common (or complex) commands.
Scalr	Provide scaling across multiple cloud computing platforms, integrates with Chef.

Conceptual Automated Toolchain

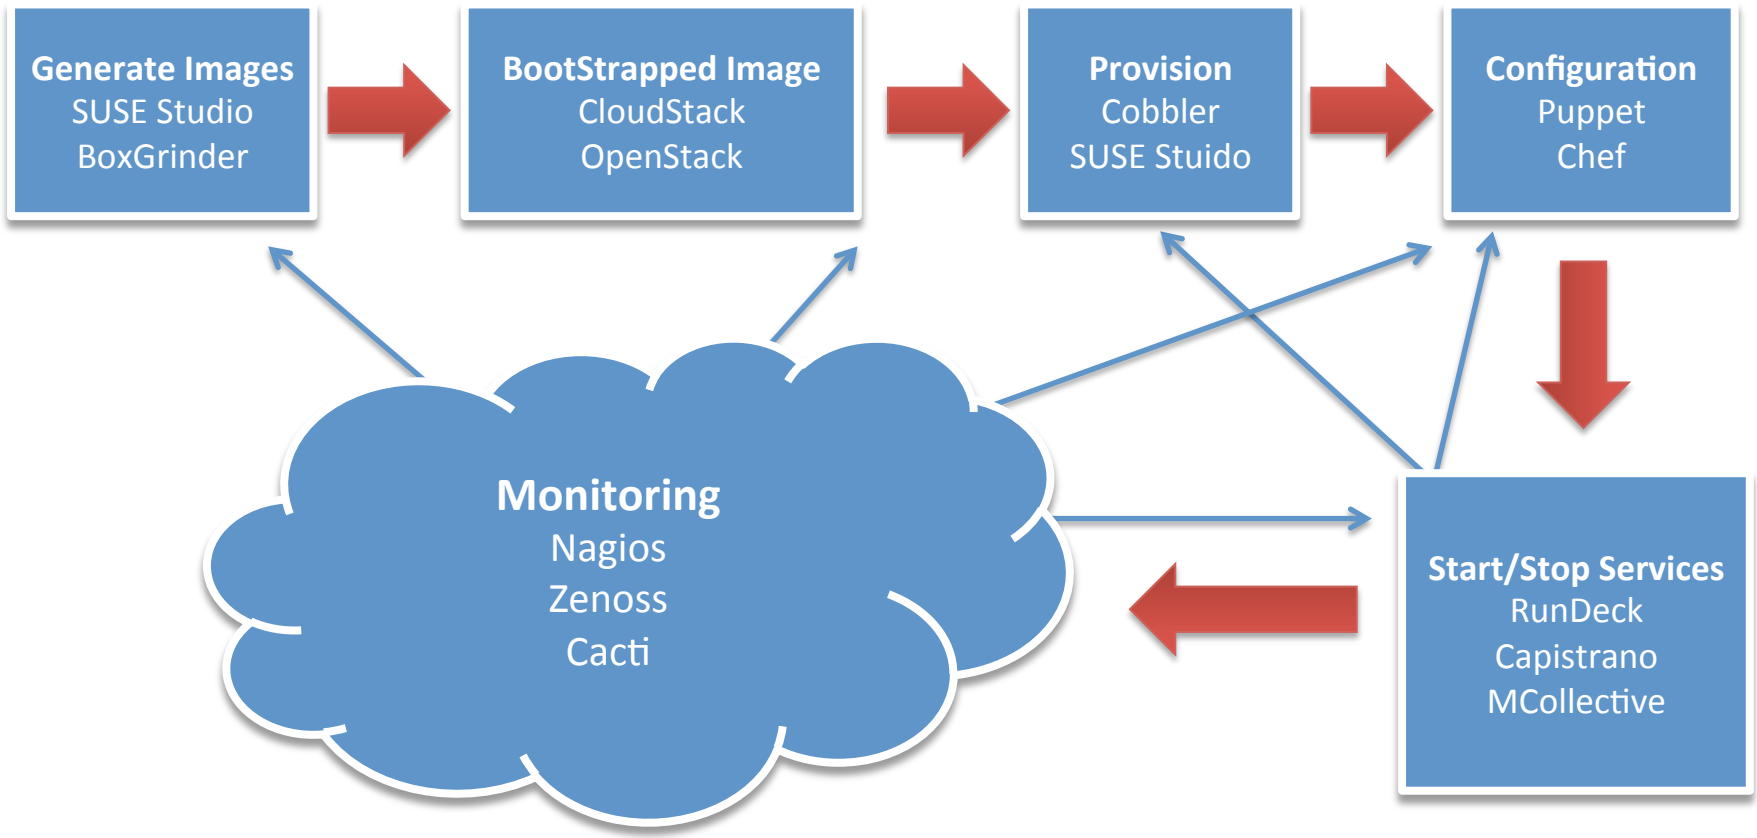

NetFlix Open Source - ToolBag for AWS

http://netflix.github.com

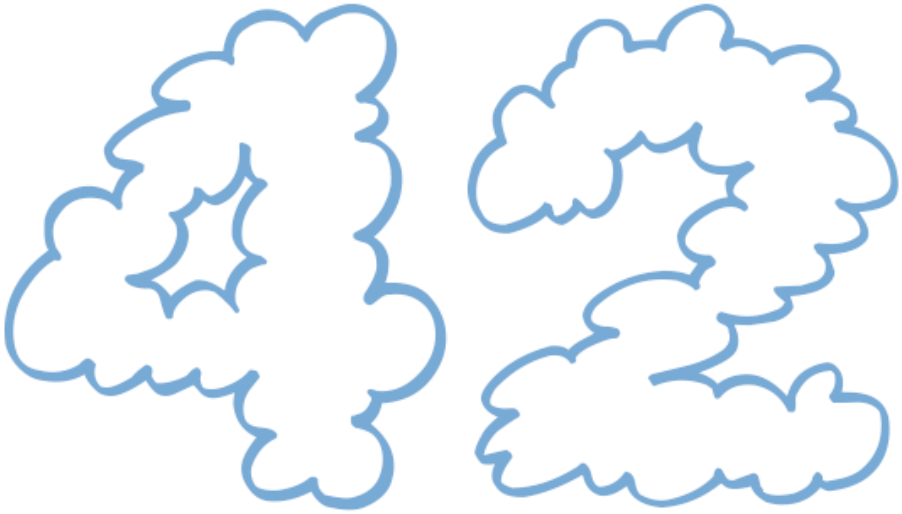

Goodbye and thanks for all the fish!

Questions?

Slides Can be Viewed and Downloaded at:
http://www.slideshare.net/socializedsoftware/

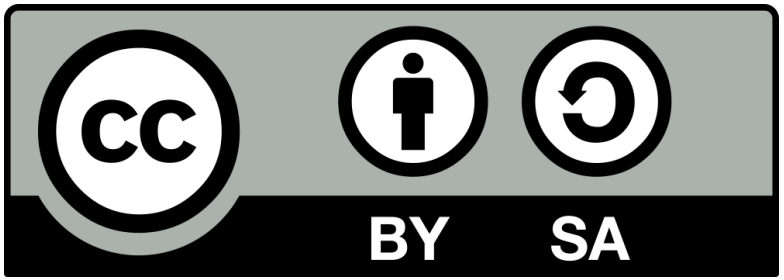

Contact Me

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

Mark R. Hinkle

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

Appendix

Additional Resources

- Devops Toolchains Group
- Software Defined Networking: The New Norm for Networks (Whitepaper)
- DevOps Wikipedia Page
- NoSQL-Database.org – Ultimate Guide to the Non-Relational Universe
- Open Cloud Initiative
- NIST Cloud Computing Platform
- Open Virtualization Format Specs
- Clouderati Twitter Account
- Planet DevOps
- Nicira Whitepaper – It's Time to Virtualize the Network
- Why Open vSwitch FAQ

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