

Open source software for building a private cloud

Michael J Pan
CEO & co-founder, nephosity

COSCUP
15 August 2010

An introduction

me

- ▶ 10+ years working on high performance (distributed, grid, cloud) computing at DreamWorks Animation, NASA JPL, NIH Center for Computational Biology, Compaq
- ▶ started nephosity in March 2010

nephosity

- ▶ develops cloud computing platform for enterprises
- ▶ showcased by STRUCTURE 2010 as one of “10 most promising cloud computing startups of 2010”

What's behind the cloud?




twitter

Gmail™
by Google

flickr™

facebook.



Motivation

Scenario

- ▶ You are (or your company is) developing a SaaS
- ▶ You require elastic compute resources

So you want to deploy in the cloud, but...

- ▶ Public clouds do not satisfy your (security, performance, etc.) requirements
- ▶ You want to use open source components in your cloud

What's available to you?

Why not Amazon EC2 (or some other public cloud?)

EC2 (more specifically, dynamic provisioning¹ capabilities provided by EC2) is only one part of the equation

- ▶ Core is dynamic provisioning capabilities
- ▶ EC2 is not open source.

You need a machine image to run on EC2– what software (OS + platform) to install on the image? What are the (open source) alternatives for dynamic provisioning?

¹the ability to start up and tear-down compute resources on-demand

What about Hadoop?

Hadoop is also only part of the equation

- ▶ Hadoop-core provides map-reduce functionality
- ▶ HDFS provides data management functionality

How do you control Hadoop jobs? What alternatives to Hadoop are there?

Cloud computing stack

- ▶ Infrastructure
 - ▶ Hypervisor / machine image
 - ▶ Dynamic provisioning
 - ▶ Operating system
- ▶ Platform
 - ▶ Data management
 - ▶ Map-reduce
 - ▶ Workflow management
 - ▶ Messaging
- ▶ Cluster management
 - ▶ Configuration
 - ▶ Analytics

- ▶ Will discuss only open source offerings that have been **released**
- ▶ Will present what's available, not how to adopt/implement them
- ▶ Lists may be incomplete
- ▶ You will see some badly hand drawn graphics



- ▶ Hypervisor / Virtual machine
- ▶ Dynamic provisioning
- ▶ Operating system

Hypervisor / Virtual machine

- ▶ Hardware virtualization
- ▶ Allows multiple virtual machines to run on a single physical machine

Hypervisor / Virtual machine

- ▶ QEMU (virtualizer)
- ▶ KVM
- ▶ Xen
- ▶ VirtualBox (Desktop only)

Dynamic provisioning

de/allocate compute resources on demand

- ▶ You get compute resources when you want them
- ▶ Compute resources are reclaimed when you release them

Open source software


- ▶ Eucalyptus
- ▶ OpenNebula / Haizea
- ▶ Condor (via VM universe)
- ▶ TCloud Elaster (not yet released)

Operating system

- ▶ The interface between your software and the underlying hardware
- ▶ In cloud computing, operating systems are stored as machine images
- ▶ Images are distributed to local storage on-demand
- ▶ Loaded into memory and booted into the hypervisor by the dynamic provisioner

- ▶ Various Linux distributions
 - ▶ Ubuntu
 - ▶ SUSE
 - ▶ Fedora
 - ▶ CentOS
- ▶ BSD (on VirtualBox)

Platform



- ▶ Data management
- ▶ Map reduce
- ▶ Workflow management
- ▶ Messaging

Data management

- ▶ Distribute your data across your network
- ▶ Replicate your data across your network
- ▶ Optimize retrieval to improve computation time
- ▶ Optimize storage requirements


Data management considerations

- ▶ SQL vs. NoSQL
- ▶ Replication degree
- ▶ small file vs. BLOB storage
- ▶ Consistency
- ▶ Centralized vs. decentralized
- ▶ Access patterns

Data management

- ▶ HDFS (Hadoop)
- ▶ SphereFS (UIC)
- ▶ DDFS (Nokia)
- ▶ Cassandra (Facebook / Apache)
- ▶ MongoDB
- ▶ CouchDB (Apache)
- ▶ MySQL (Oracle)
- ▶ PostgreSQL
- ▶ Ceph (DreamHost), release as part of Linux v2.6.34

Map reduce




- ▶ Split and parallelize a task into many parts
- ▶ Combine the results of the split tasks for a final result

Open source offerings

- ▶ Hadoop (Yahoo)
- ▶ Sphere (UIC)
- ▶ Disco (Nokia)

Workflow management



- ▶ design
- ▶ specification
- ▶ coordinated execution

of compute tasks

Open source offerings

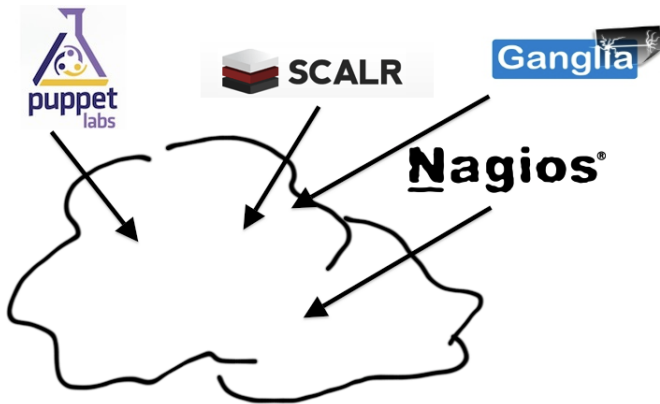
- ▶ Oozie (Yahoo)
- ▶ Pig (Hadoop / Apache)
- ▶ Cascading (Concurrent)
- ▶ Azkaban (LinkedIn)
- ▶ pomsets (nephosity)

- ▶ Unified framework for your application and all components to communicate with each other
- ▶ Above the network hardware and network protocol layer
- ▶ Your application handles only discrete messages

Open source offerings

- ▶ qpid (Apache)
- ▶ RabbitMQ (SpringSource / VMWare)
- ▶ ZeroMQ (iMatix)

Cluster management



Configuration management

- ▶ Configuration of your running cloud instances
- ▶ Software upgrades
- ▶ Dynamic configuration that cannot be stored onto OS images
- ▶ Relaxes storage constraints vs. using OS images

Open source offerings

- ▶ Chef (Opscode)
- ▶ Puppet
- ▶ StarCluster (MIT)



- ▶ Collection and visualization of the status of your cloud
 - ▶ Compute load
 - ▶ Network usage
- ▶ Dynamic load balancing and scaling of your cloud
 - ▶ Start new instances
 - ▶ Tear down existing instances

Open source offerings

- ▶ Graphite (Orbitz)
- ▶ Scalr
- ▶ Nagios
- ▶ Ganglia

Questions?

For more info:

Michael Pan

mjpan@nephosity.com