Exploring the Cloud

Ezra Zygmuntowicz http://engineyard.com



What is Cloud Computing?

What is Cloud Computing?

laaS

What is Cloud Computing?

laaS

SaaS

What is Cloud Computing?

laaS

SaaS

PaaS

What is Cloud Computing?



HaaS

SaaS

PaaS

What is Cloud Computing? HaaS laaS SaaS PaaS

The Definition has been Clouded...

If you listen to the pundits, anything on the internet is now 'cloud computing'

I'm Going to Define Cloud Computing as:

Virtualized Compute, Memory and Storage

API accessible

Resources scale up/down on demand

AWS is the prototypical example

So what does it really take to work with these low level cloud resources?

A Different Way of Thinking

It's all about the Automation

Three Pillars of Automation

Repeatable

Idempotent

Treat an instance like a referentially transparent function

f(config) -> Fully Configured Instance

Calling f(config) will *always* create the same instance whenever config == config

Treat instances as throw away

Prefer rebuilding from scratch

Only persist what truly needs to be persistent

Puppet

CFEngine

Soon to be released Mystery Project ;)

2. Ad Hoc Change

Run this command on this server now

Deploy this new code to these servers now

Gather information from these servers now

2. Ad Hoc Change

Vertebra

Capistrano

2. Monitoring

You can't fix what you can't see

2. Monitoring

Nagios Munin Collectd Ganglia Monit



Engine Yard abstracted from Engine Yard



Clouds are too low level for average humans



First Target: AVVS



Custom Gentoo Linux State based config management Ad hoc Change **Extensive Monitoring** 24/7 support High Level Cloud

Bridge multiple providers Highly tuned databases on EY Elastic app servers on ec2 Will support any compelling new cloud platforms

Vertebra: p2p Cloud Control

Problem:



Problem:

Booting a cluster is a complex process



Problem: Lots of interdependent tasks to perform



Problem:

No 'init' for clusters...



Problem:

Lots of servers to control in the cloud



Problem: Complex interactions and deployment scenarios



Solution: Vertebra



Operations on Resource Sets is Vertebra's Fundamental Abstraction

<op token='7e5293c579f4a9fa8da6320e9ef03e3f'
 type='install'>
 <res>/cluster/1</res>
 <res>/slice/42</res>
 <res>/gem</res>
 <string name='gem'>hpricot</string>
 </op>

Resources are Hierarchical

Needs /cluster/l /slice/42 /gem

Needs /cluster/l /slice/42 /gem

Provides /cluster/l /slice/42 /gem

Needs /cluster/l /slice/42 /gem



Provides /cluster/l /slice/42 /gem

Needs /cluster/ /slice/42 /gem

SI Also Matches! Provides /cluster/l /slice/42 /gem

(slice 42 on *all* clusters)

Needs /cluster/l /slice/ /gem

SI Also Matches! Provides /cluster/l /slice/42 /gem

(all slices on cluster I)

Needs /cluster/l /slice/42 /gem

No Match!

Provides /cluster/l /slice/42 /xen

(no soup for you)

Resource Discovery

Who provides this set of resources?

Resource Discovery

Who provides this set of resources?

>> @agent.discover('/cluster/ey04', '/slice/42')
=> {"jids"=>["ey04-s00042@ey04.engineyard.com/agent"]}

Resource Discovery

Who provides this set of resources?

>> @agent.discover('/cluster/ey04', '/slice/42')
=> {"jids"=>["ey04-s00042@ey04.engineyard.com/agent"]}

Dynamic DNS for Cloud Resources









Fault Tolerant Protocol on top of XMPP

> Each agent can act as a Client or a Server or both at once.

State machines for Client/Server Protocols

Fault Tolerant Protocol on top of XMPP



Key/Value Data Store

Queried via an 'op'

Uses Resource 'Sets' as compound keys

key: /cluster/l /customer/foobar

value: {... arbitrary hash ...}

Workflow coordination of multiple agents

Move Slice Workflow

input: /cluster/1, /slice/42, /node/2



Operations are like unix processes: hash as stdin/stdout

Workflows are like command line pipes tying together multiple operations

Integration with IM via XMPP

O O foo@conference.ey04.engineyard.com			\bigcirc
👻 🖳 🥥 😃 🥯 🗞 🎯 🌖 A			主 🖭 📀
Ezra Zygmuntowicz %help	2:32		<pre> ey04-s00015 ey04-s00014 ey04-s00013 ey04-s00013 ey04-s00012 </pre>
I am ey04-s00013@ey04.engineyard.com/agent I provide these resources: <actor: 13,="" provides="/cluster/ey04," reload="" slice=""> <actor: 13,="" gem="" provides="/cluster/ey04," slice=""></actor:></actor:>			<pre> ey04-s00011 ez </pre>
ey04-s00012	2:32		
I am ey04-s00012@ey04.engineyard.com/agent I provide these resources: <actor: 12,="" provides="/cluster/ey04," reload="" slice=""> <actor: 12,="" gem="" provides="/cluster/ey04," slice=""></actor:></actor:>			
ey04-s00014	2:32		
I am ey04-s00014@ey04.engineyard.com/agent I provide these resources: <actor: 14,="" provides="/cluster/ey04," reload="" slice=""> <actor: 14,="" gem="" provides="/cluster/ey04," slice=""></actor:></actor:>			
ey04-s00011	2:32		
I am ey04-s00011@ey04.engineyard.com/agent I provide these resources: <actor: 11,="" provides="/cluster/ey04," reload="" slice=""> <actor: 11,="" gem="" provides="/cluster/ey04," slice=""></actor:></actor:>			
ey04-s00015	2:32		
I am ey04-s00015@ey04.engineyard.com/agent I provide these resources: <actor: 15,="" provides="/cluster/ey04," reload="" slice=""> <actor: 15,="" gem="" provides="/cluster/ey04," slice=""></actor:></actor:>			

Any 'state' in a workflow can call out to meatware

If we have time...

