





# Scaling Ruby from the Inside Out

Ezra Zygmuntowicz





# Does Rails Scale?



# Does Rails Scale?

Of course it does...



# Does Ruby Scale?



# Does Ruby Scale?

Of course it does...



Should you use Rails or  
Merb?



Should you use Rails or  
Merb?

YES!



Are you sick and tired  
of endlessly debating  
this crap?



Are you sick and tired  
of endlessly debating  
this crap?

Of course you are...



Me too, let's talk about  
something more  
interesting...



# ~~Scaling Ruby from the Inside Out~~



What should we talk  
about then?



What should we talk  
about then?

The weather maybe?



Forecast calls for 100%  
chance of clouds...



# Cloud Computing





# Cloud Computing Platforms

Abstract/Virtual Resources

CPU

Memory

Storage

API Access



# Cloudy Weather

- Amazon EC2
- Engine Yard
- Joyent
- Mosso
- 3terra
- Google App Engine
- Project Caroline
- GoGrid

This is only the beginning...



There will be \*many\*  
cloud computing  
platforms in the near  
future



Many virtual machines  
to manage



SSH doesn't scale...



SSH doesn't scale...

Takes time to setup/teardown  
connections



# SSH doesn't scale...

Takes time to setup/teardown  
connections

Inherently push only technology



Capistrano is an awesome tool



Capistrano is an awesome tool

but running commands on  
thousands of machines in  
parallel is pushing it further  
then it wants to go



How will we manage  
the complexity?



How will we manage  
the complexity?

It's all about the  
distribution...





This background compliments of the unofficial Scarface site: Make Way for The Bad Guy: SCARFACE (1983)  
Copyright © 1983 Universal City Studios, Inc. All rights reserved.

<http://www.jgeoff.com/scarface/>

We've got to think big now!  
New York, Chicago, LA



# Introducing Vertebra

Next Generation Cloud  
Computing/Automation  
Framework



# Enter XMPP

Extensible Messaging and Presence Protocol



# Enter XMPP

Extensible Messaging and Presence Protocol

## Extensible

I like that



# Enter XMPP

Extensible Messaging and Presence Protocol

## Messaging

We want to send messages to our thousands of machines



# Enter XMPP

Extensible Messaging and Presence Protocol

## Presence

We want to know when our machines are online or not  
and act accordingly



XMPP is a realtime messaging protocol built for IM/chat, great for communication between thousands of people, why not machines?



# Summoning Slices via XMPP

The screenshot shows an XMPP chat window titled "foo@conference.ey04.engineyard.com". The chat history includes a message from Ezra Zygmuntowicz with a profile picture and the text "%help". Below this, five messages from agents "ey04-s00013" through "ey04-s00015" are shown, each with a timestamp of "2:32". Each agent message follows a similar format: "I am ey04-sXXXX@ey04.engineyard.com/agent", "I provide these resources:", and two lines of XML resource information: "<Actor: provides=/cluster/ey04, /slice/XX, /reload>" and "<Actor: provides=/cluster/ey04, /slice/XX, /gem>". On the right side of the chat window, a contact list is visible, listing "ey04-s00015", "ey04-s00014", "ey04-s00013", "ey04-s00012", "ey04-s00011", and "ez". The bottom of the window shows the current contact name "foo@conference.ey04.engineyard.com" in a text input field.

foo@conference.ey04.engineyard.com

Ezra Zygmuntowicz 2:32  
%help

ey04-s00013 2:32  
I am ey04-s00013@ey04.engineyard.com/agent  
I provide these resources:  
<Actor: provides=/cluster/ey04, /slice/13, /reload>  
<Actor: provides=/cluster/ey04, /slice/13, /gem>

ey04-s00012 2:32  
I am ey04-s00012@ey04.engineyard.com/agent  
I provide these resources:  
<Actor: provides=/cluster/ey04, /slice/12, /reload>  
<Actor: provides=/cluster/ey04, /slice/12, /gem>

ey04-s00014 2:32  
I am ey04-s00014@ey04.engineyard.com/agent  
I provide these resources:  
<Actor: provides=/cluster/ey04, /slice/14, /reload>  
<Actor: provides=/cluster/ey04, /slice/14, /gem>

ey04-s00011 2:32  
I am ey04-s00011@ey04.engineyard.com/agent  
I provide these resources:  
<Actor: provides=/cluster/ey04, /slice/11, /reload>  
<Actor: provides=/cluster/ey04, /slice/11, /gem>

ey04-s00015 2:32  
I am ey04-s00015@ey04.engineyard.com/agent  
I provide these resources:  
<Actor: provides=/cluster/ey04, /slice/15, /reload>  
<Actor: provides=/cluster/ey04, /slice/15, /gem>

ey04-s00015  
ey04-s00014  
ey04-s00013  
ey04-s00012  
ey04-s00011  
ez

foo@conference.ey04.engineyard.com



# Ejabberd

Horizontally scalable/federated erlang XMPP/Jabber server

High performance even with tens of thousands of connections

Forms the backbone of the system



Ruby based agents run on  
every slice/instance/node  
in the cloud



Presence let's us know when  
agents come online or go  
offline



Security/Discovery/  
Introduction Framework  
Keeps Agents in Touch with  
Each Other.



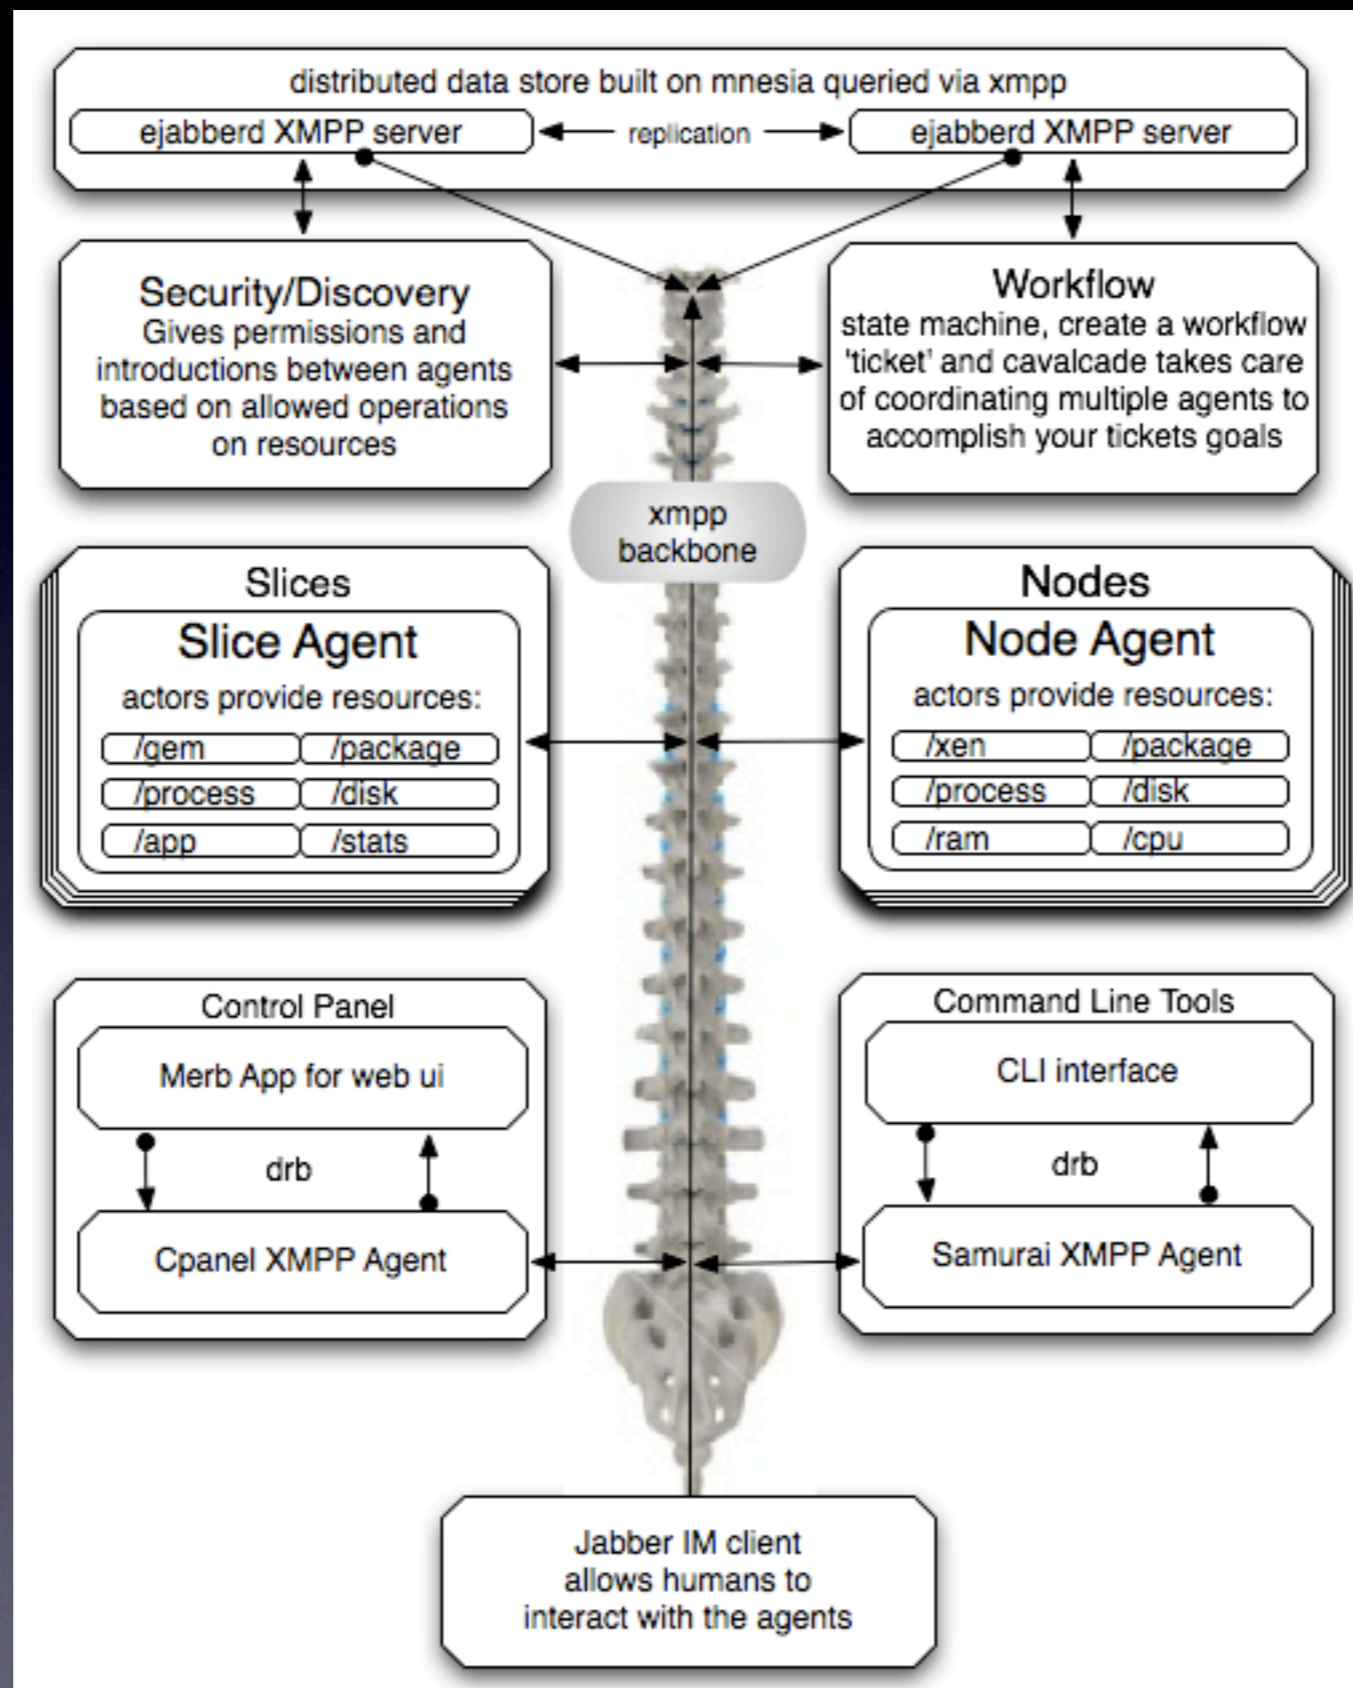
# Distributed Hierarchical Key- Value Storage System Queried via XMPP



State Machine/Workflow  
Engine Coordinates Multiple  
Agents/Humans



# Vertebra





CREST



CREST

Cloud-REST



# CREST

## Cloud-REST

Operations on Resources  
in the cloud



operation:list  
/cluster/ey05  
/node/42  
/xen

```
module Vertebra  
  module Actors  
    class Xen < ::Vertebra::Actor  
      provides '/xen'  
      register_with_agent  
  
      def list(options = {})  
        Pars::Xm::ListOutput.new(`xm list`).parse  
      end  
    end  
  end  
end  
end
```



# Resource Discovery



# 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"]}
```




Map/Reduce  
Scatter/Gather Dispatch



# Map/Reduce

## Scatter/Gather Dispatch


Discover






# Map/Reduce Scatter/Gather Dispatch

Discover




Scatter






# Map/Reduce Scatter/Gather Dispatch

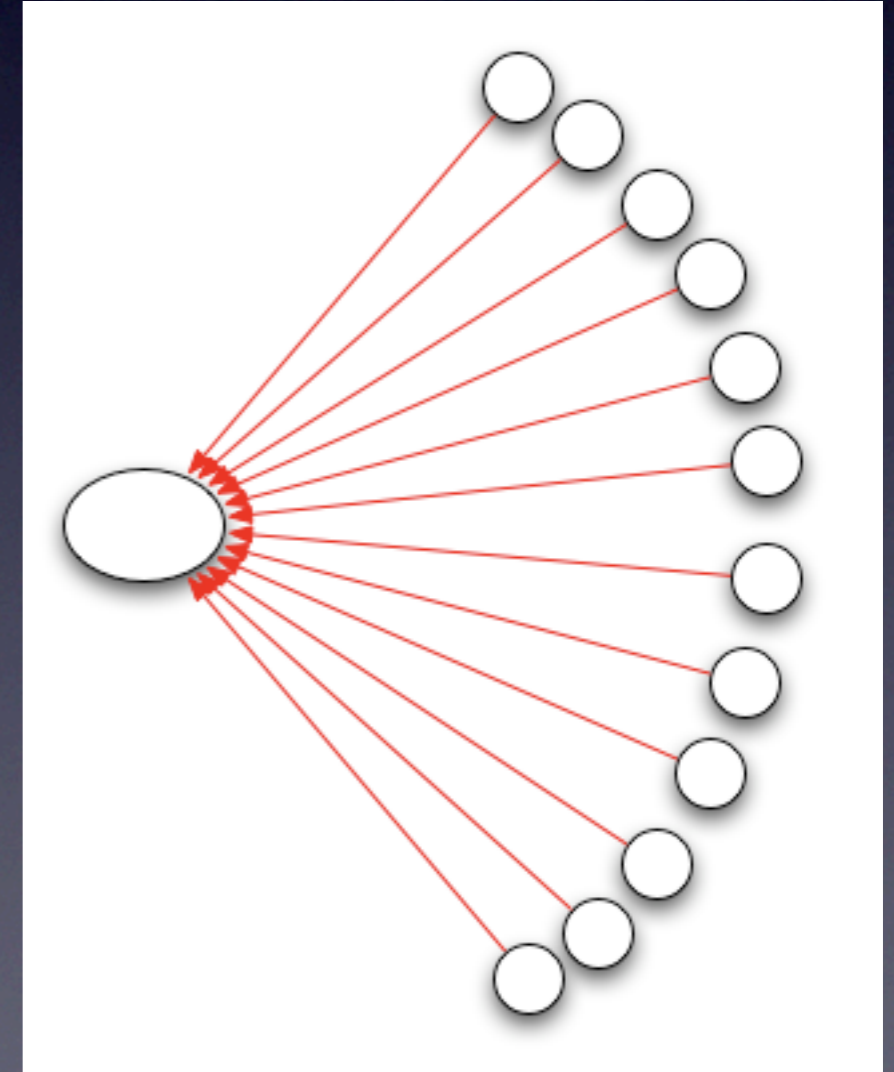
Discover



Scatter



Gather





# quick demo

```
>> @agent.broadcast('list', '/cluster/ey04', '/node/', '/xen')
# outputs:
Node: ey04-n01
Slice: Domain-0, cpus: 8, memory: 1024
Slice: ey04-s00008, cpus: 2, memory: 4096
Slice: ey04-s00011, cpus: 2, memory: 4096
Slice: ey04-s00041, cpus: 1, memory: 640

Node: ey04-n01
Slice: Domain-0, cpus: 8, memory: 1024
Slice: ey04-s00028, cpus: 4, memory: 8192
Slice: ey04-s00034, cpus: 1, memory: 640
Slice: ey04-s00038, cpus: 1, memory: 640
...
```



Not only server  
automation



Ejabberd + Comet + Jsjac

==

Realtime Web Platform



Many Ideas to  
Experiment With



XMPP > HTTP



Watch this space...



One more thing...



One more thing...

Introducing:  
Engine Yard Express

[http://express.engineyard.com/system/ey\\_express\\_v0.1.tar.bz2](http://express.engineyard.com/system/ey_express_v0.1.tar.bz2)



# Engine Yard Express

GNU GRUB version 0.97 (638K lower / 653248K upper memory)

Engine Yard Express v0.1

Use the ↑ and ↓ keys to select which entry is highlighted.  
Press enter to boot the selected OS, 'e' to edit the  
commands before booting, or 'c' for a command-line.

The highlighted entry will be booted automatically in 2 seconds.



# Express

gentoo linux™

[http://express.engineyard.com/system/ey\\_express\\_v0.1.tar.bz2](http://express.engineyard.com/system/ey_express_v0.1.tar.bz2)



Questions?