

Cloud... so much more than a tools fest




Qcon London 2012

Patrick Debois



Vagrant &
Veewee



<http://jedi.be/blog>




@patrickdebois

CUTTER
CONSORTIUM
●●● ACCESS TO THE EXPERTS

www.cutter.com

 Atlassian

Freelancer




Context:
a traditional
enterprise







With a
developer group
on a mission
to ***innovate***



Realtime TV - Web interaction





Second Screen Applications

ENTERPRISE CAT



**IS READY FOR
PROVISIONING**

Maximum warp!

TV-show
+/- 1 million
concurrent viewers



ENTERPRISE CAT



IS RE THIS FOR
PROV IS NOT A SIGN IING

Cloud Cat

Doin' his cloud thing

PALAS®

CAT'S EYE
GLASS MARBLES

INTERNATIONAL DISTRIBUTORS .

Google App
Engine




A close-up photograph showing a large number of small, grey tabby kittens. They are being held and supported by a person's hands, which are visible at the bottom of the frame. The kittens are packed together, filling most of the image. Some kittens are looking towards the camera, while others are looking away. The background is slightly out of focus, showing what appears to be a dark-colored car wheel and a person's leg in dark clothing. The overall scene is very cute and heartwarming.

Abstraction is AWESOME

If it fails
hard to debug/
understand



Too generic




Java Rails Esper Redis Hadoop ETL tools
Python Nodejs MongoDB Mysql

Need more control



**Oh, here's your
problem ...**

IAAS



The computer is missing

cat


I never did any
cloud before.
Only enterprise
stuff

i is not




What I learned from managing Production Servers

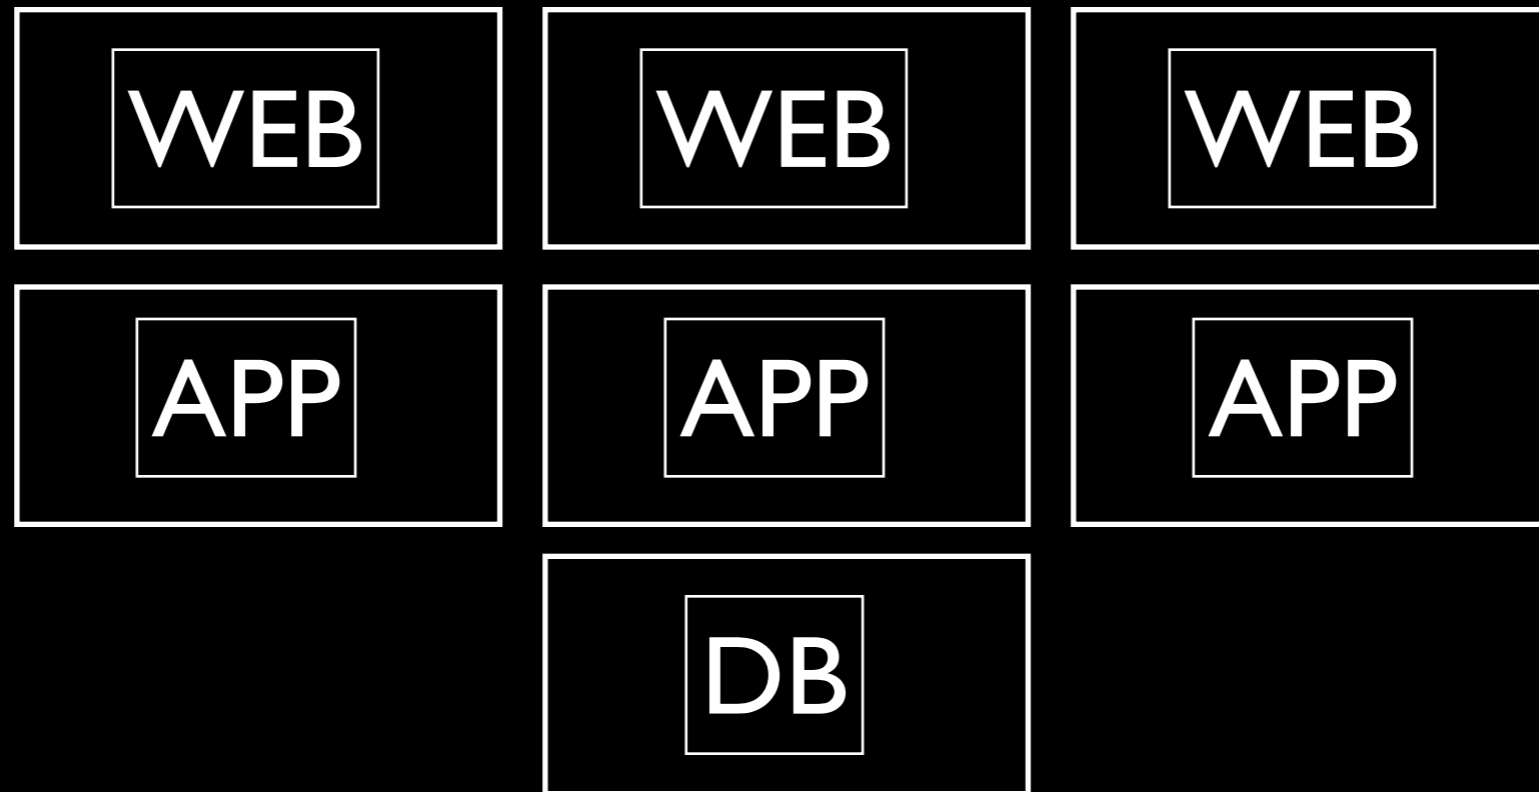
A Single Server




Starts Growing



Generic Servers become Specific Servers

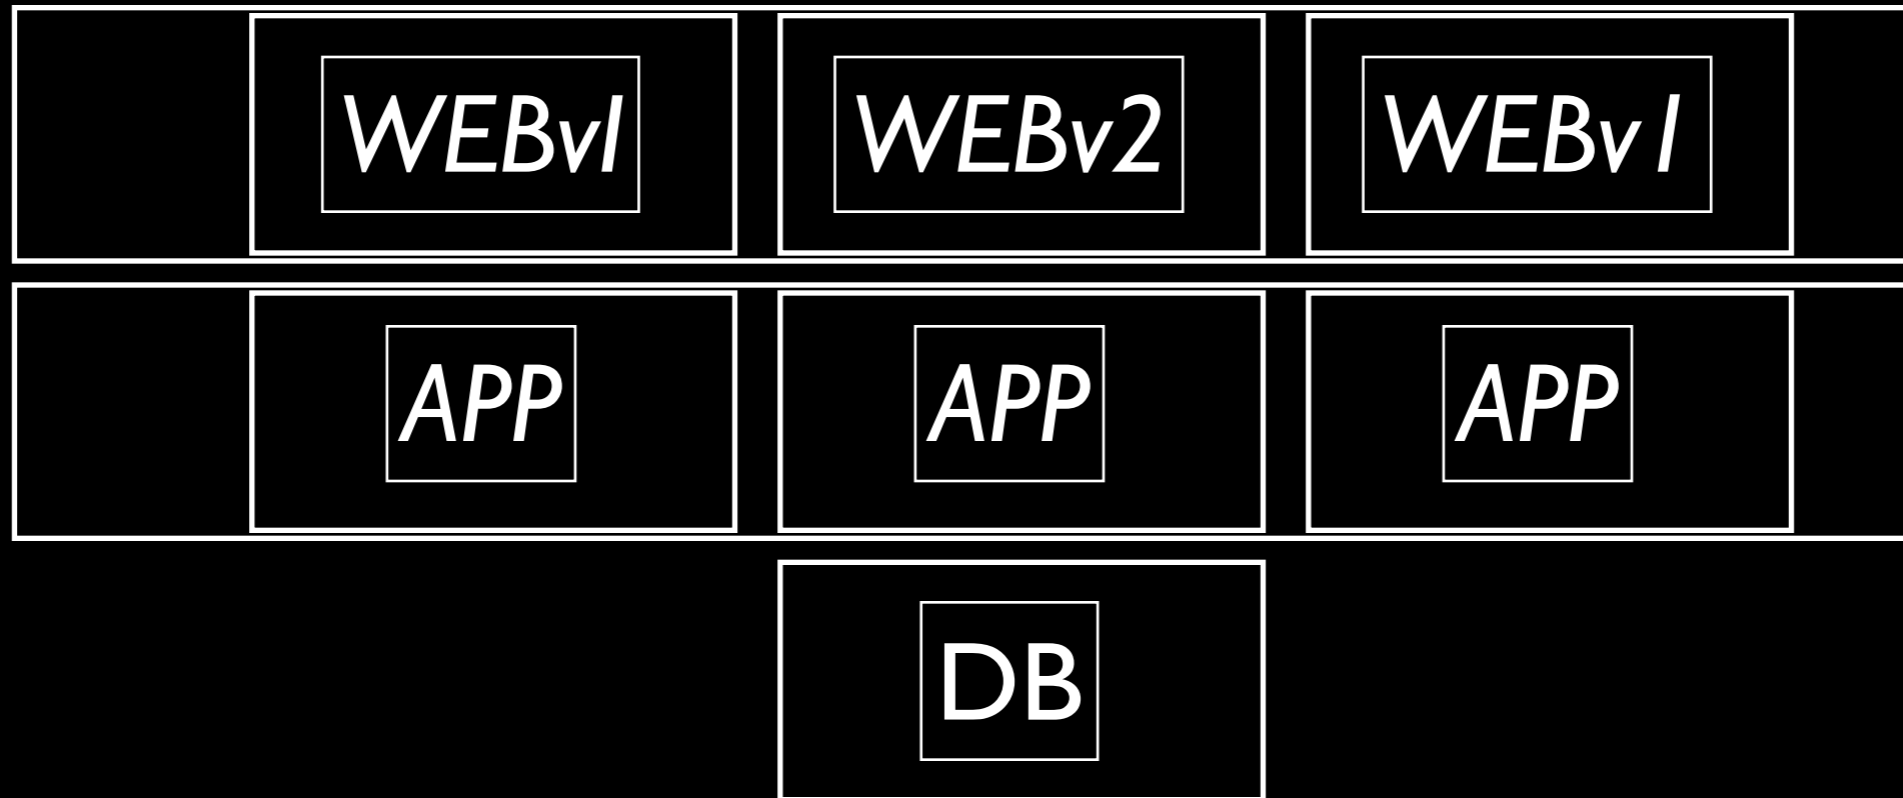


Physical becomes Virtual

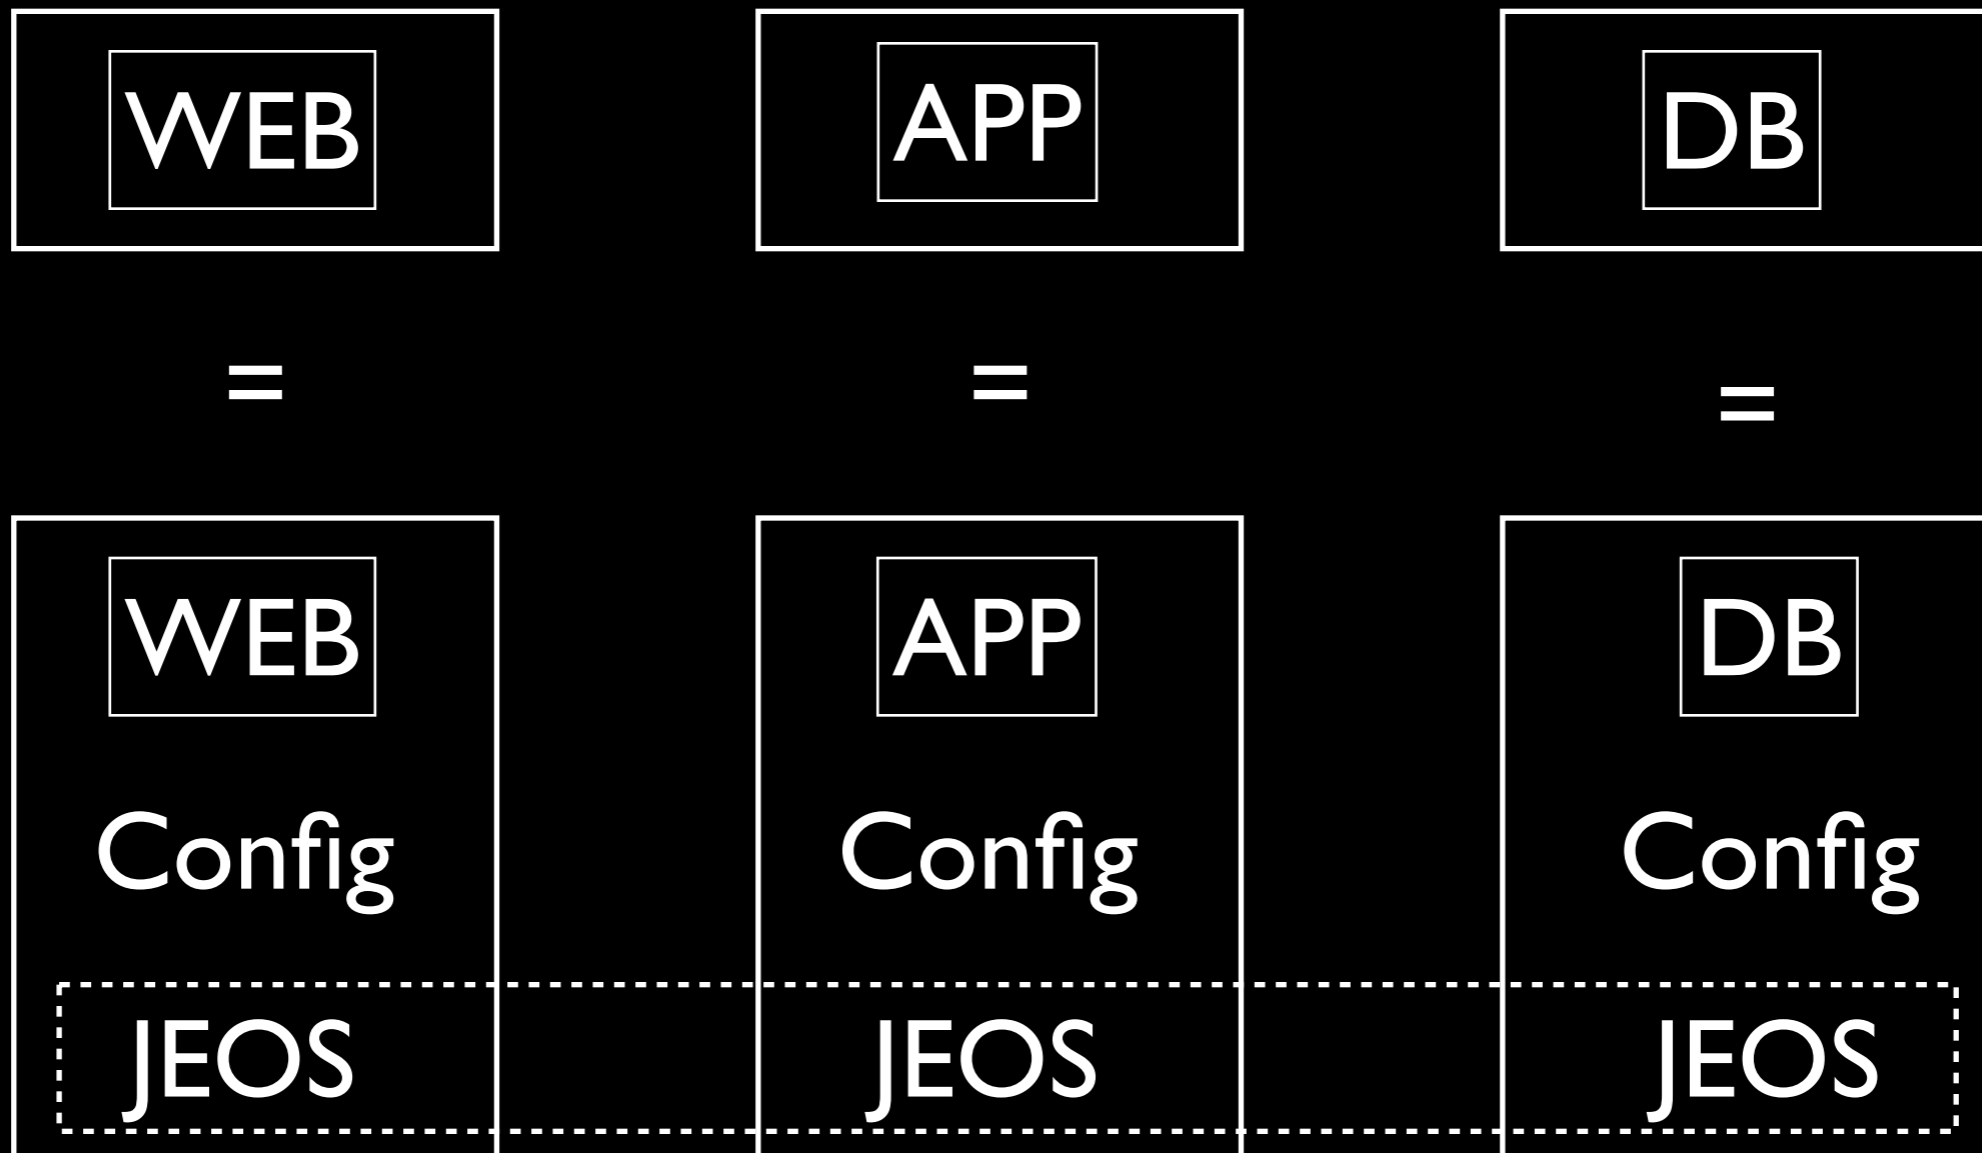


xen/vsphere/kvm/...

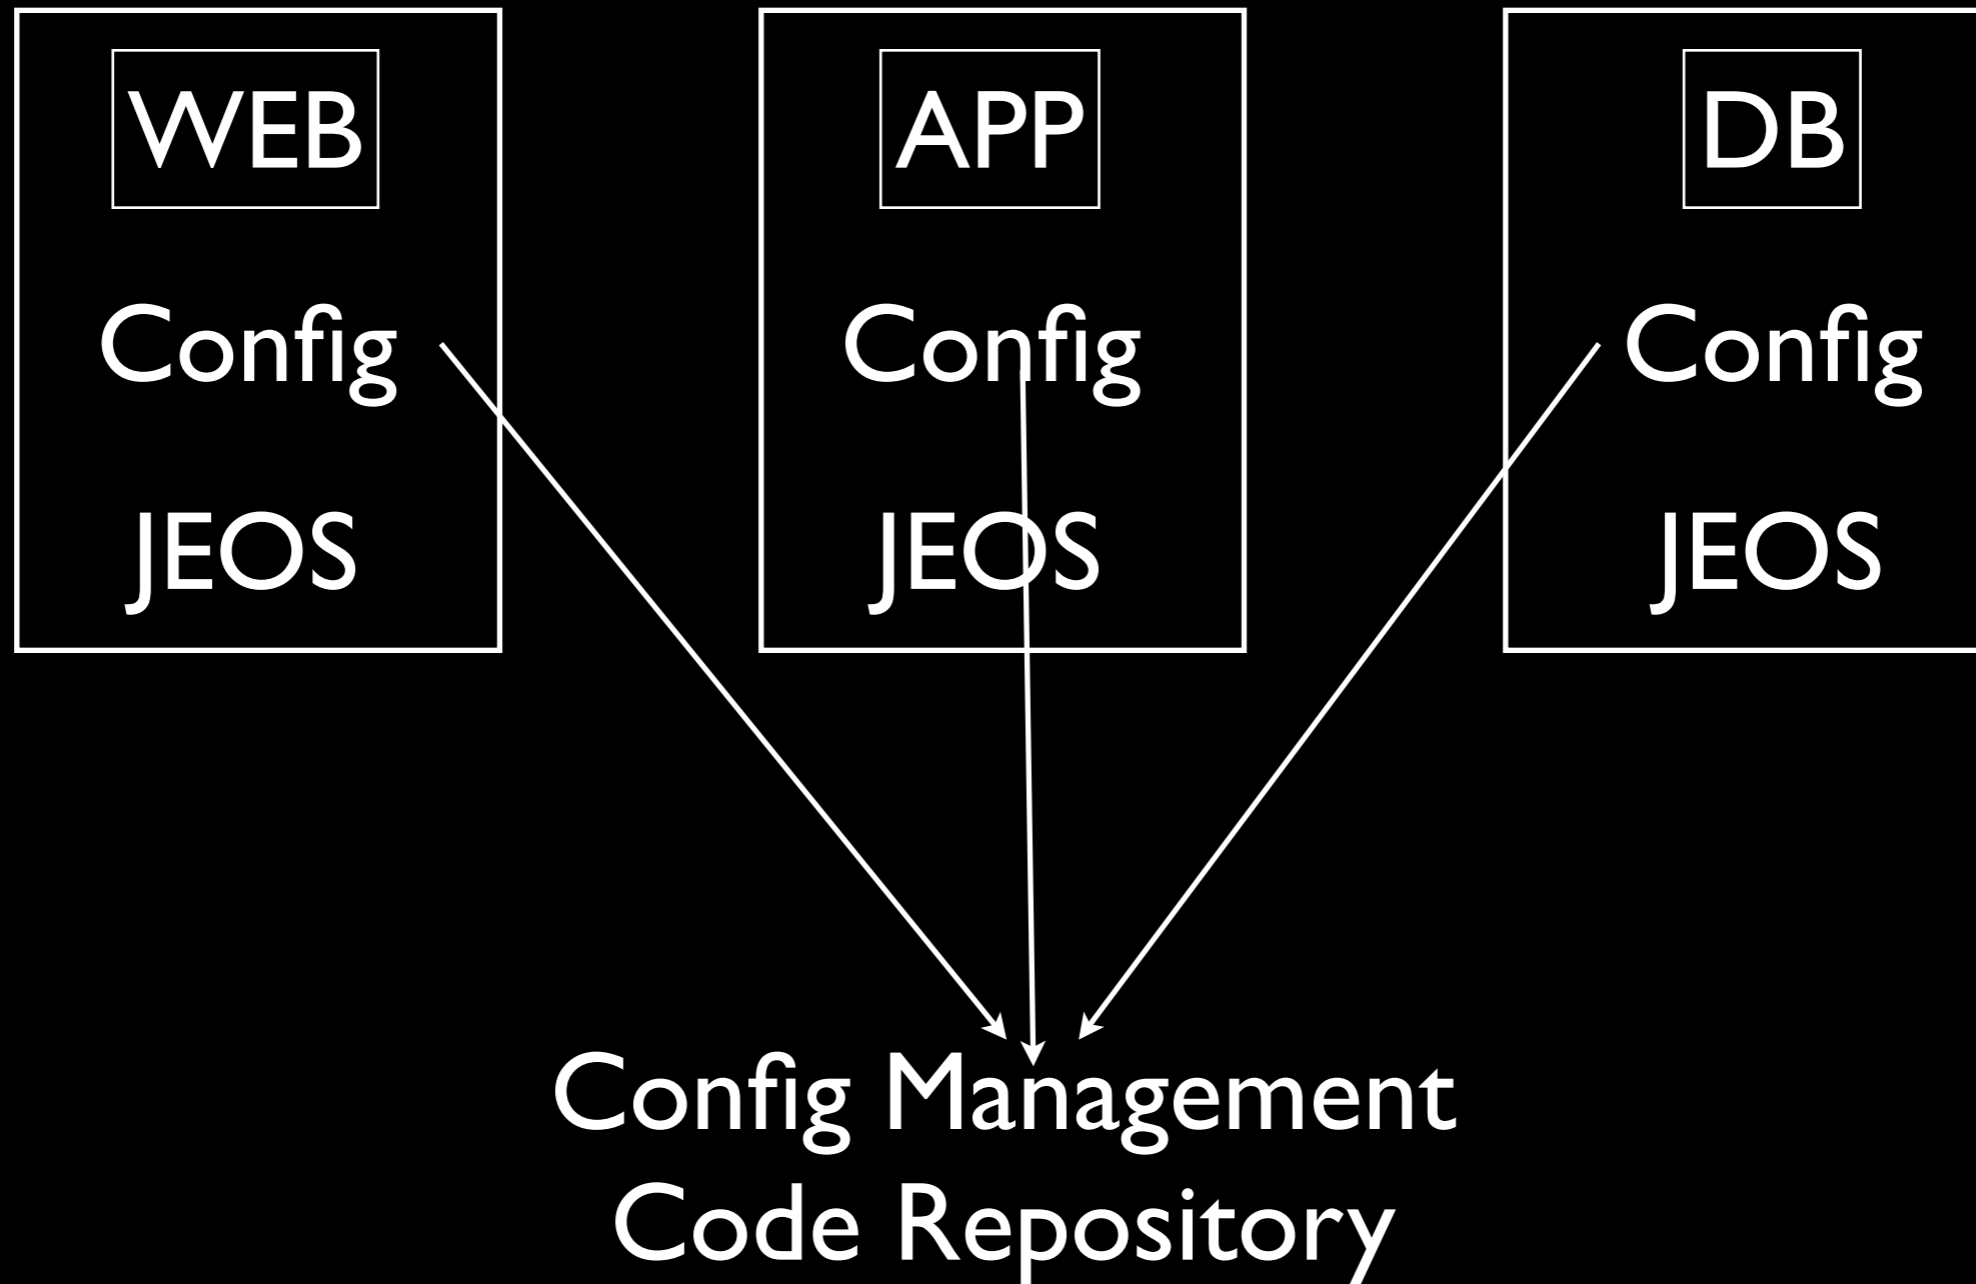
We learned cloning isn't working



We introduce config management



Infrastructure as code



Cfengine/Puppet/Chef

What I learned from managing Test Servers



Reuse Across Environments

DEV

TEST

PROD

WEB

WEB

APP

APP

DB

DB


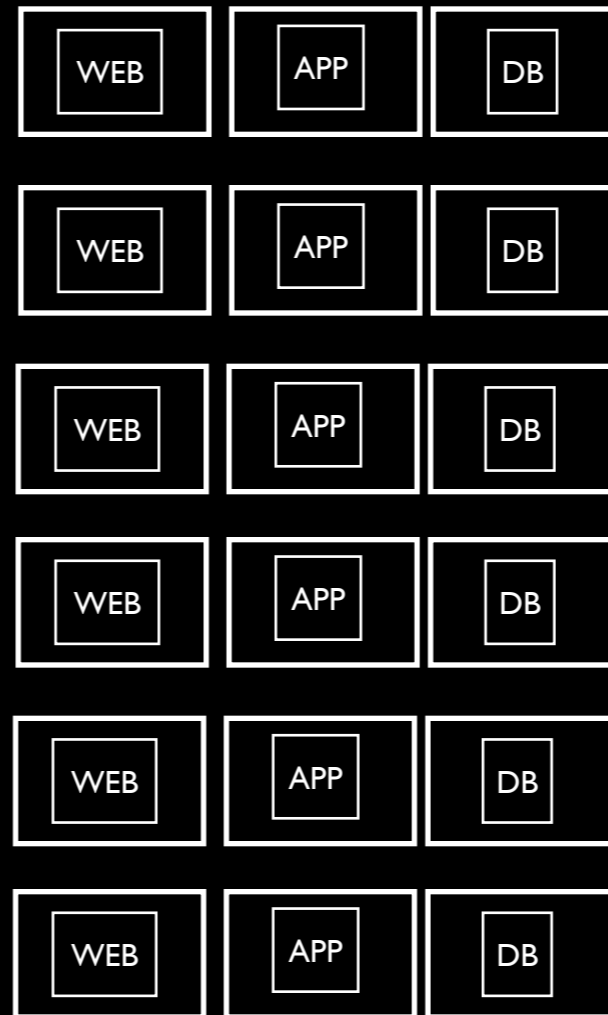
Infrastructure
Code Repository

Explosion of VM creation

DEV

TEST


PROD



Infrastructure
Code Repository


Application
Code Repository

UI Interface to Automated Provisioning of VMS



Cobbler, Spacewalk, ...

Metadata Registry for Systems



Config
Management

Puppet/Chef
Server

we hearz thingz



from de other side...



**You had beans,
didn't you**



OH LINUX?



**I'VE NEVER INSTALLED A
GUI FOR THAT.**

```
$ knife ec2 server create \  
-r 'role[webserver]' -I ami-7000f019 \  
-f m1.small \  
-A 'Your AWS Access Key ID' \  
-K 'Your AWS Secret Access Key'
```



```
$ puppet node create --image ami-XxXXxXXX \  
--keypair puppetlabs.admin --type m1.small
```

aka “cloudpack”



Servers

Security groups

Elastic Load balancers

Keys

S3 Storage
E-Mail service



<http://fog.io>
@geemus



Custom Scripts

AWS Network Constraints



Security
Groups only on
creation

Only 1
network
interface


ELB
not on non-
standard ports

Dynamic Monitoring

```
nodes = search(:node, "hostname:[*  
TO *] AND  
chef_environment:#{node.chef_enviro  
nment}")
```



VM creation failure, network hickups, disk erratic behavior




ICANHASCHEEZBURGER.COM 🍷 🍷 🍷




Embrace Fail.

Architect cat sez



Re-architect

bak to drawing bord for joo



“Quis custodiet
ipsos Custodes”

who watches the watchers



Noops

What I learned working in the (Amazon)cloud



what's the difference ?

Internal

Web,App,DB

Config Mgmt

Metadata Registry

JEOS

VM

UI Provision

Cloud

Web,App,DB

Config Mgmt

Metadata Registry

AMI

Xen

AWS Console

From console **to API**

Web UI

AWS Console



“Internals” API

AWS API



Abstracted
API

Fog
Jclouds
Boto

Beyond Servers Components

Server(s)

+

Loadbalancers

DNS Service

IP Address

Email Service

EBS Volume

Firewall

Keys

From server to **stack**


Cloudformation

Json file specifying
order of component creation
and dependencies

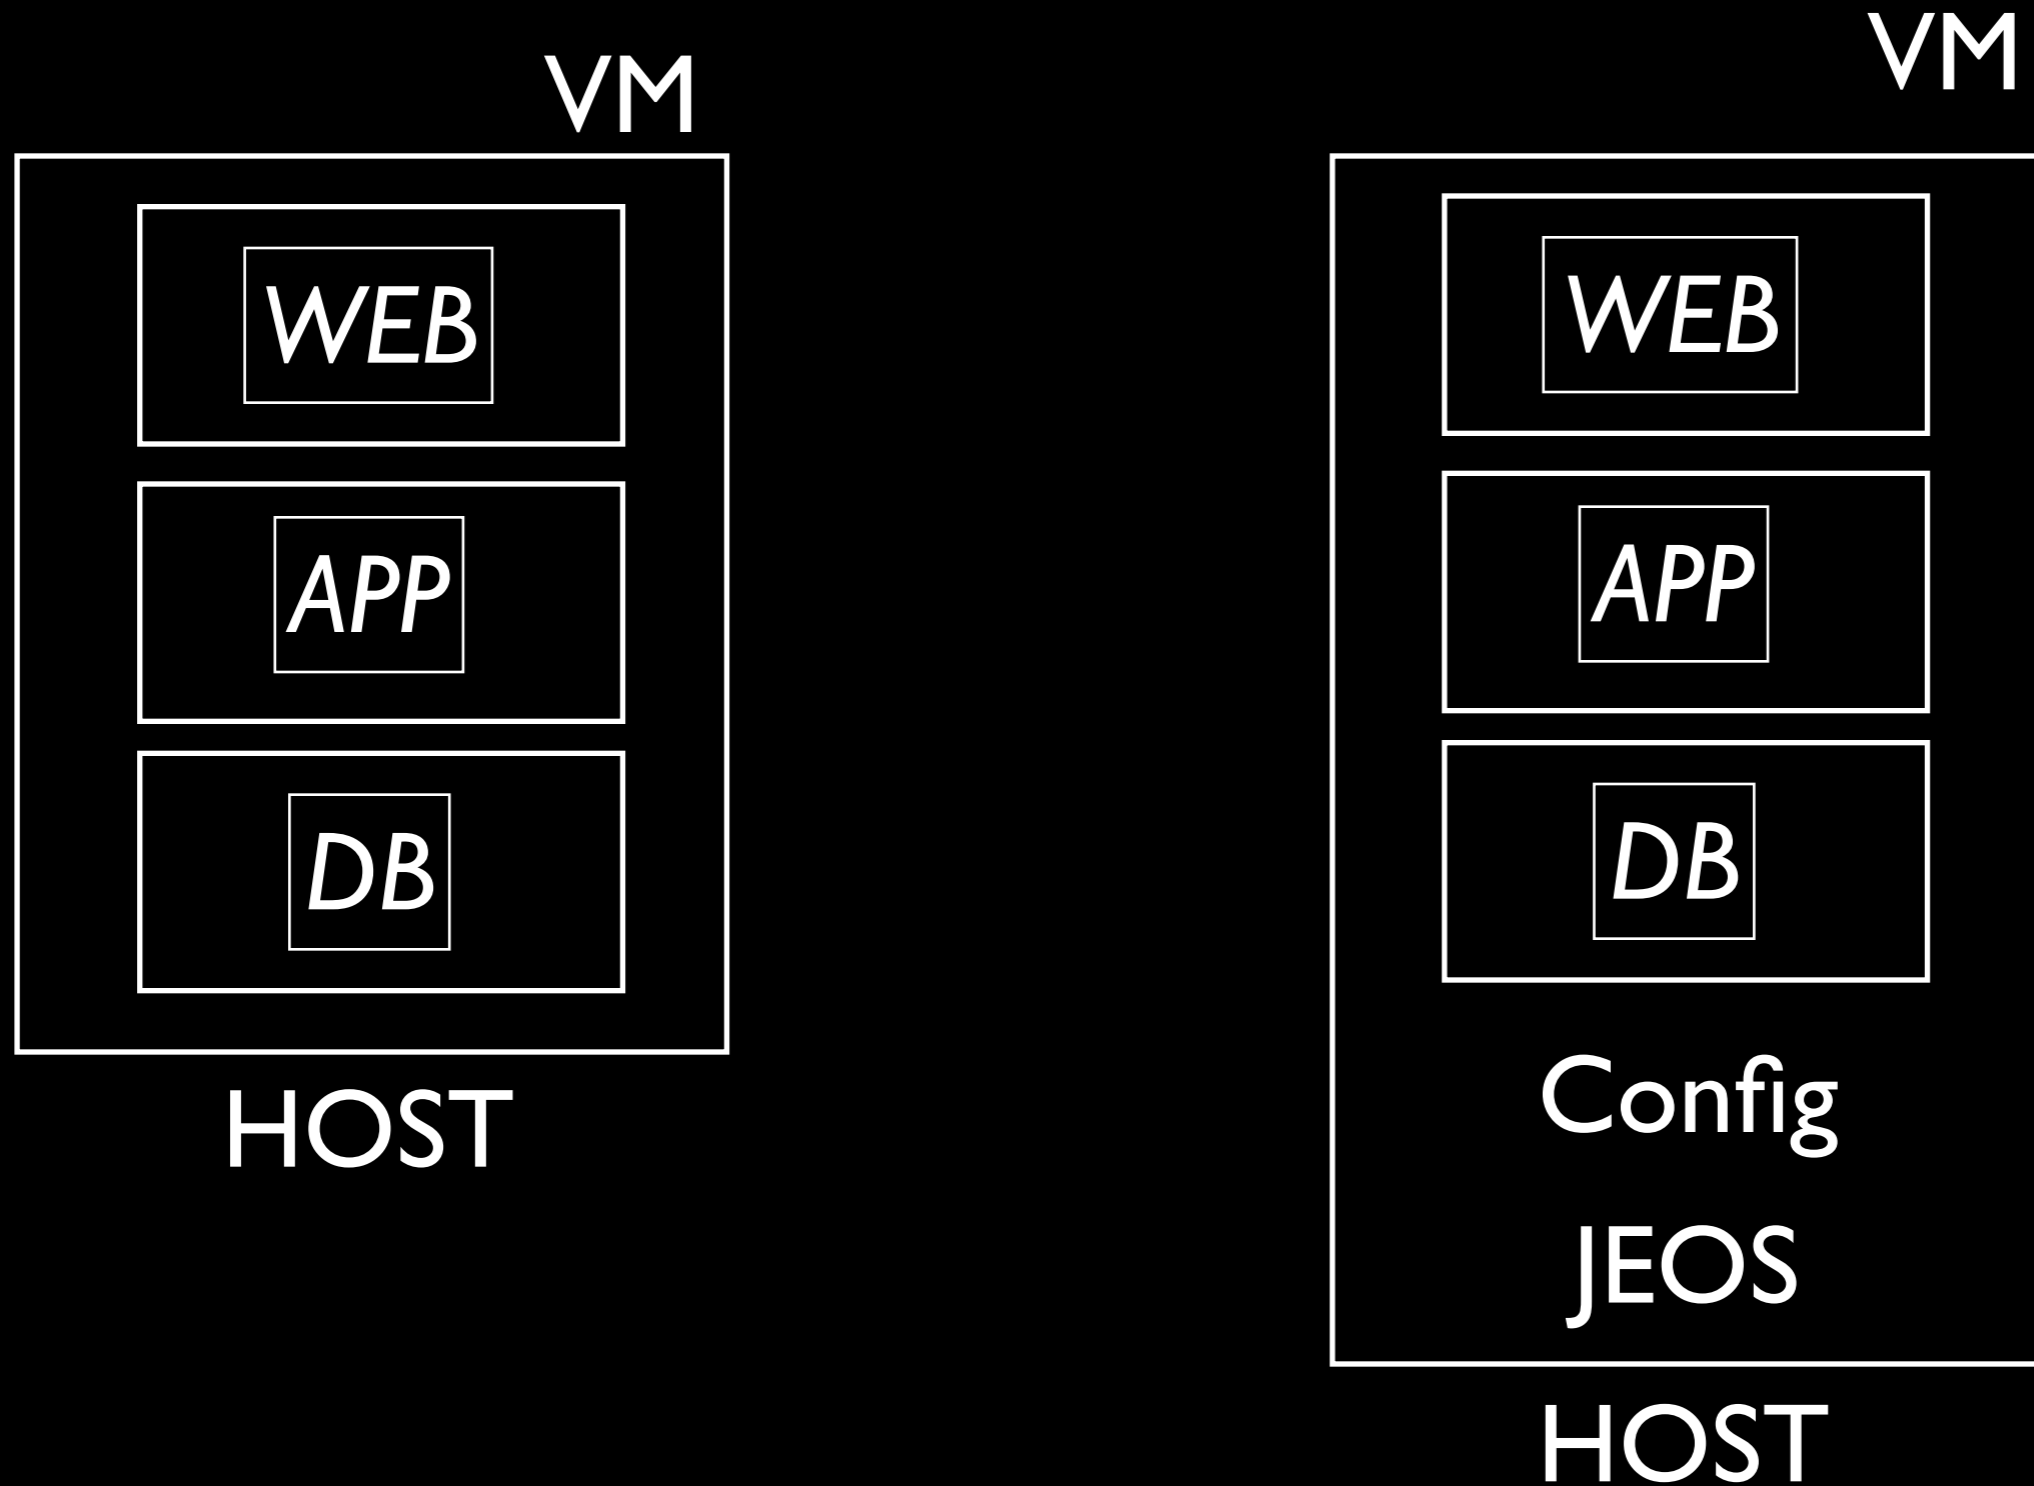


What I learned from managing development Servers

Development moves
from host into **virtual machines**



Development
starts using **config mgt**



Reuse “code” across Environments

DEV

TEST

PROD

WEB

WEB

WEB

APP

APP

APP

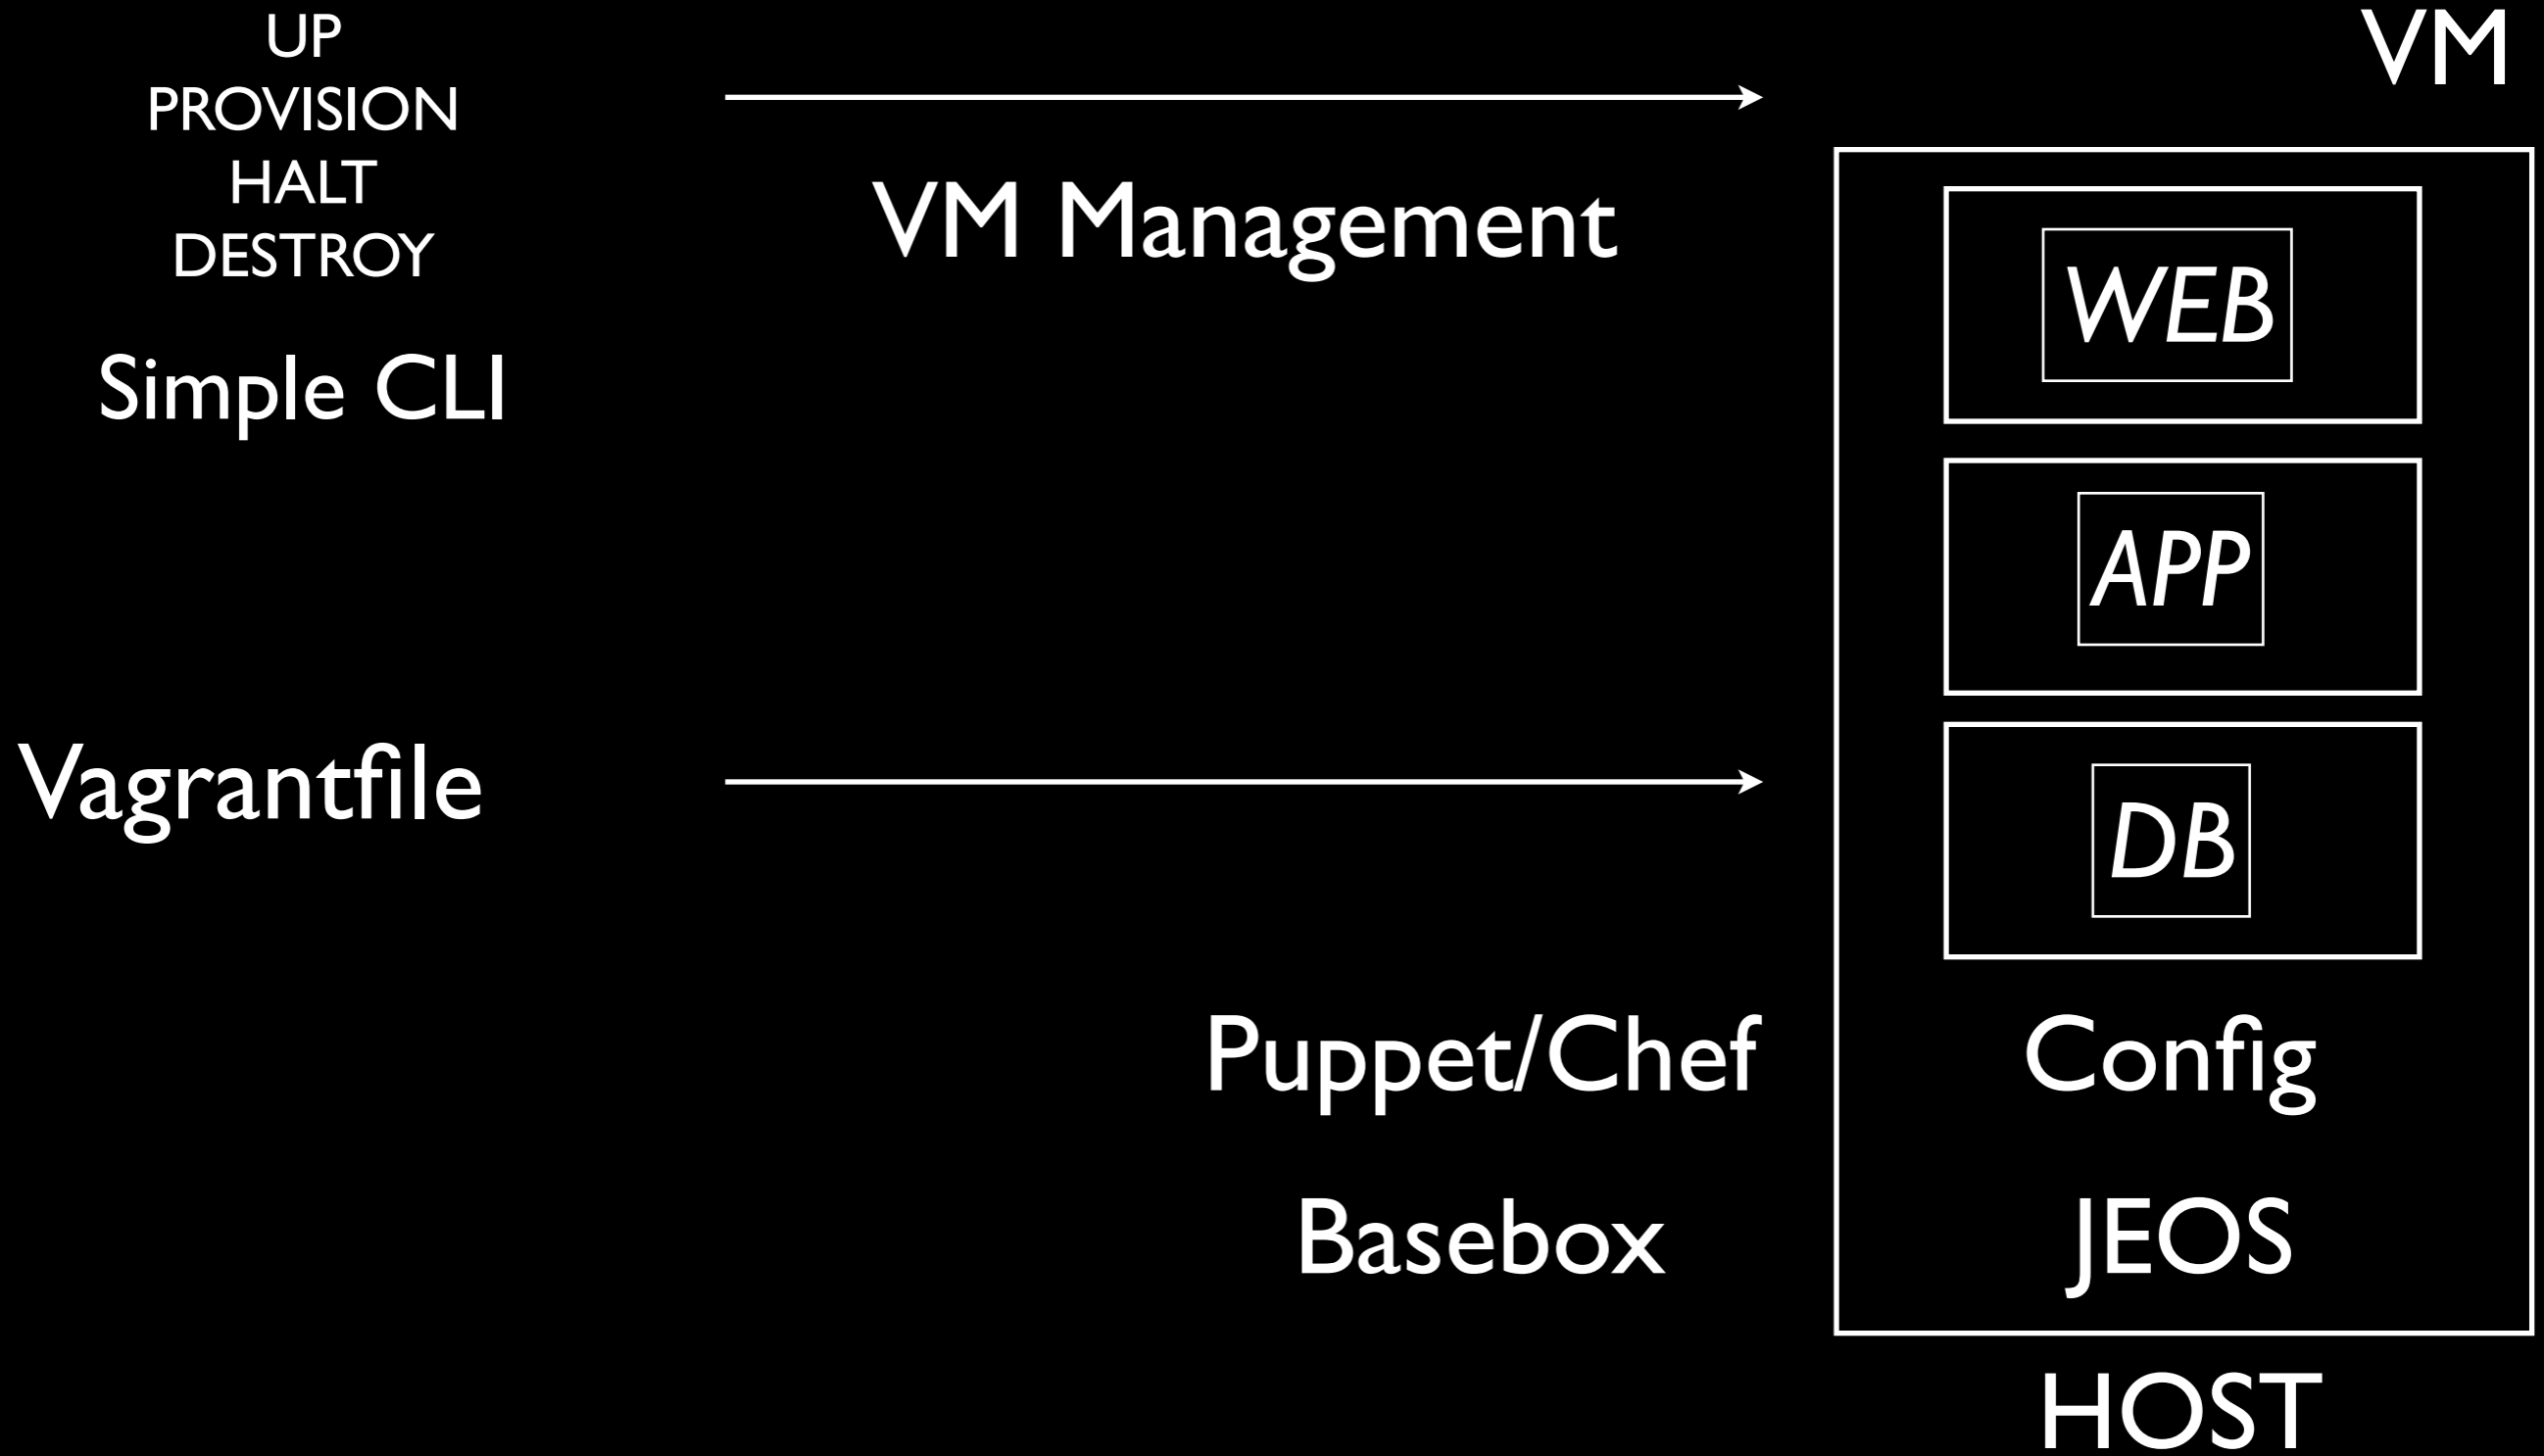
DB

DB

DB

Infrastructure
Code Repository

Vagrant



Integrate with **Continuous Integration**

DEV

TEST

PROD

APP



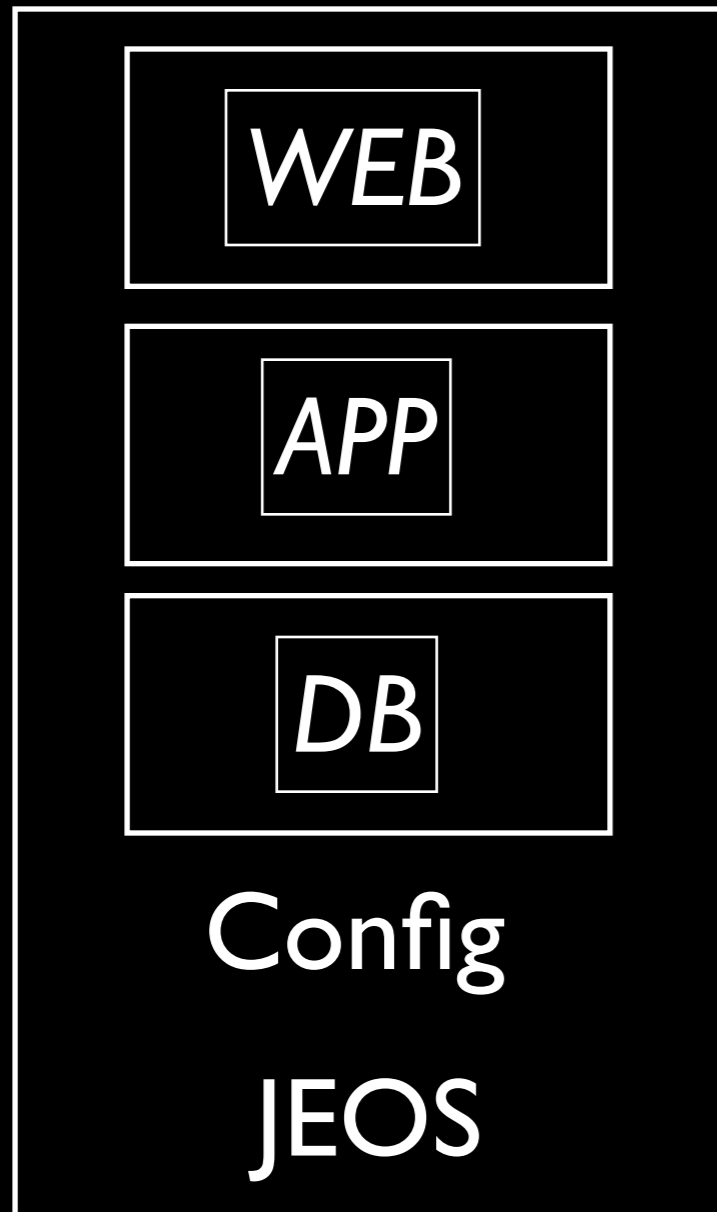
INFRA



Infrastructure
Code Repository

Application
Code Repository

Development/VM **workflow**



VM

UP

PROVISION


HALT

DESTROY


TEST



DEV



PROD



APP


INFRA

Setup Outgrew
VM on Laptop

Setup Outgrew
Test Lab

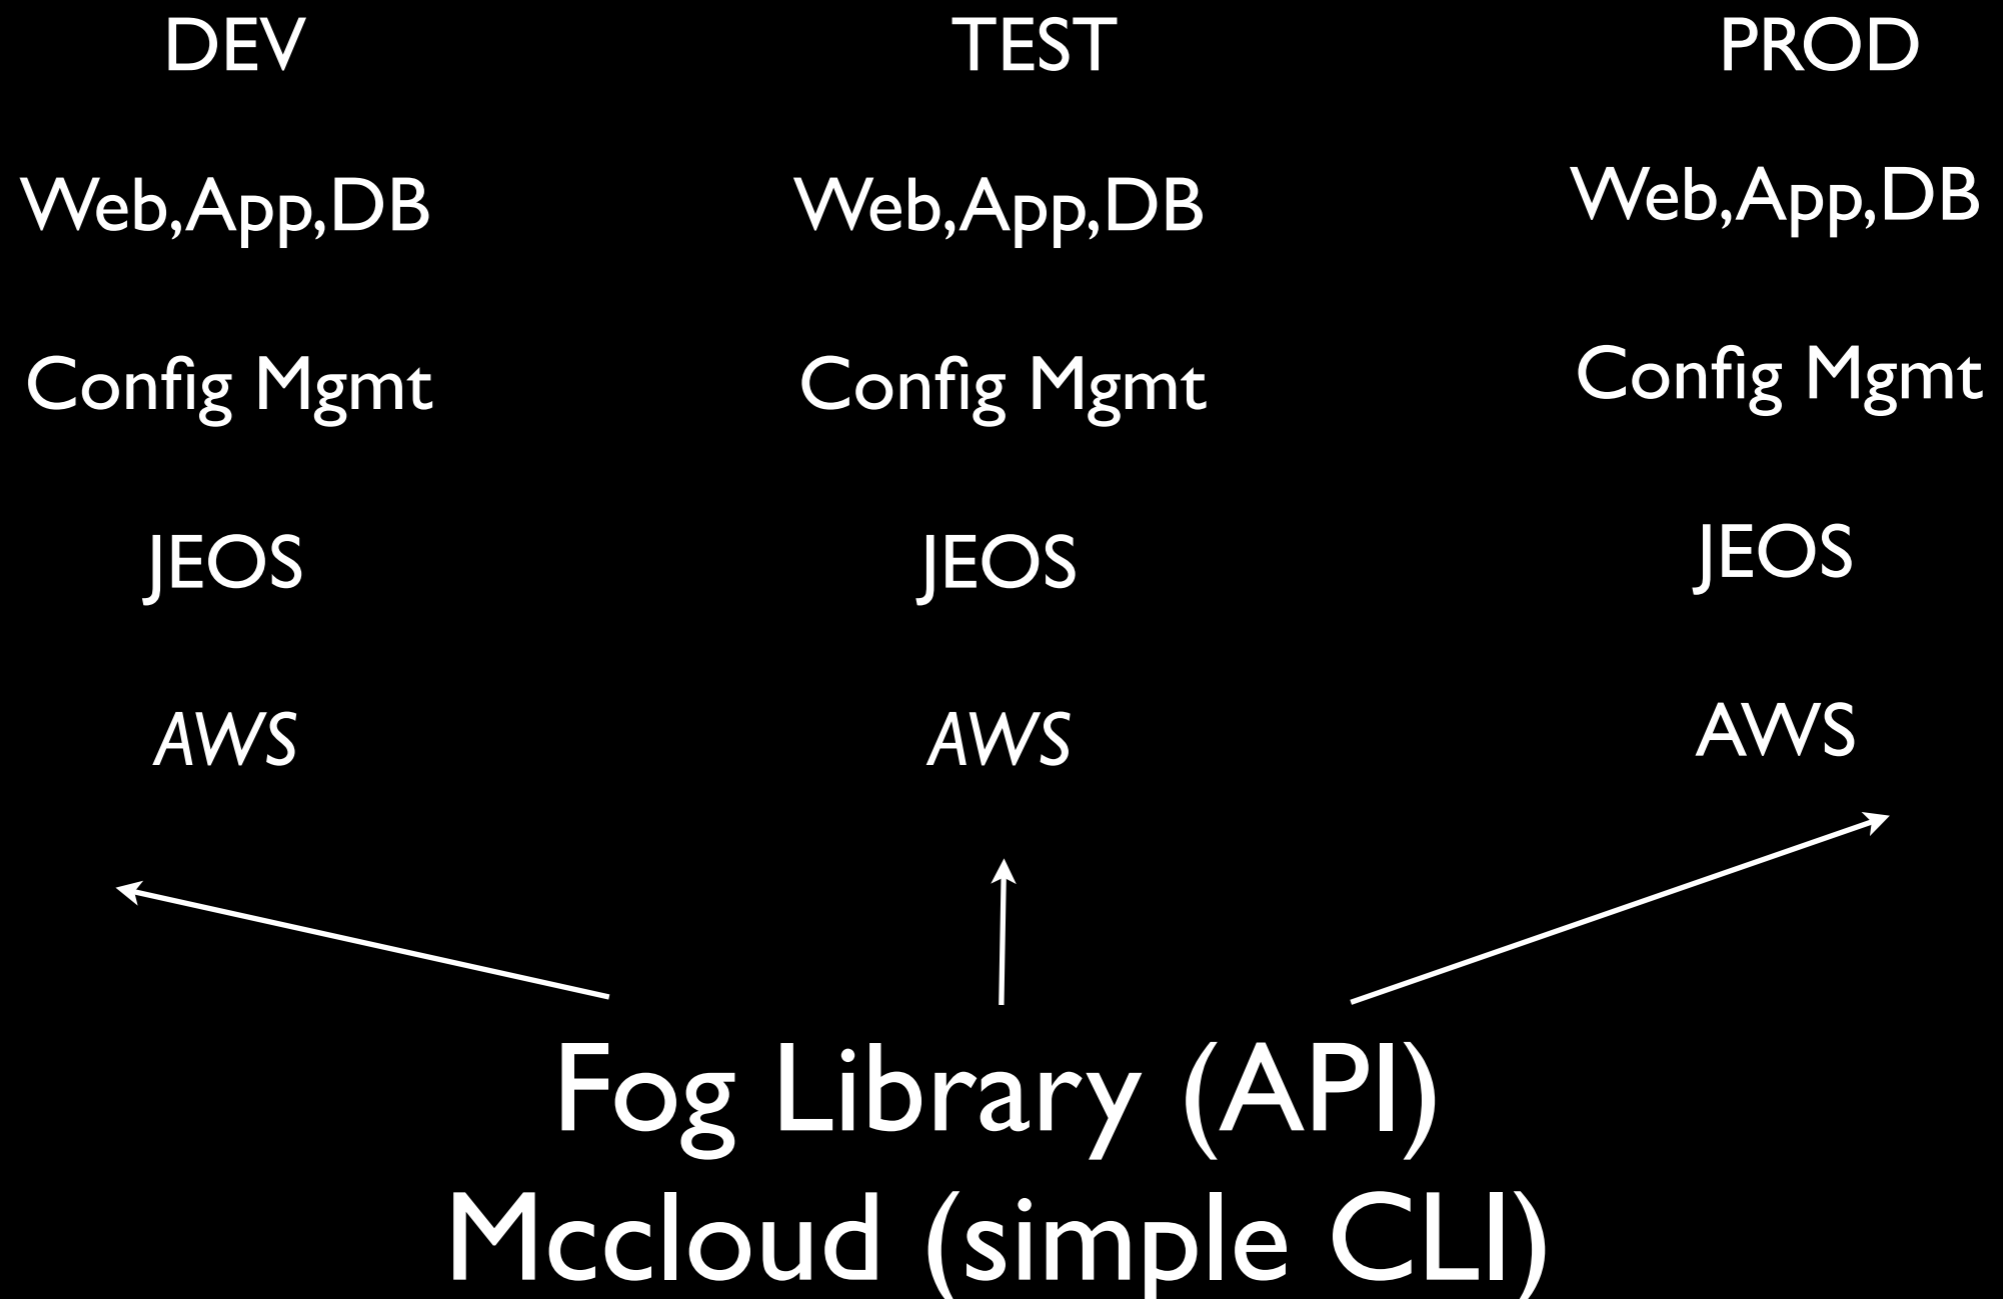
Peak Capacity
needed in Prod

CLOUD



Reuse “workflow” across Environments

“If it’s hard to it more often”



+Reuse workflow across hypervisors customers

Web,App,DB

Config Mgmt

JEOS

Virtualbox

Web,App,DB

Config Mgmt

JEOS

KVM


Web,App,DB

Config Mgmt

JEOS

AWS

Fog Library



Cloud Libs (Jclouds/Fog/Boto) embracing **old and personal**

“personal”

Virtualbox

VM Fusion

“old”

Vsphere

Kvm

Libvirt

“hybrid”

Openstack

Eucalyptus

“new”

AWS

Rackspace

Abstracting OS Installation

Debian Ubuntu Solaris Win Redhat

Archlinux Suse Centos

Kickstart Preseed

Unattended.xml

Virtualbox KVM Fusion Parallels

<http://github.com/jedi4ever/veewee>

Cloud Libs (Jclouds/Fog/Boto)

beyond servers

DNS

Keys

IPs

Storage

Securitygroups

Loadbalancer

Workflow **beyond servers**

VM

up
provision
halt
destroy

Other

ip
balance
sorry

Self Servicing



Continuous Integration to Continuous Delivery

Faster/Delivery



DEV


TEST

PROD

APP



INFRA



Faster/Feedback



Infrastructure
Code Repository

Application
Code Repository

Confidence

Faster/Delivery



DEV

OPS

Faster/Feedback



Infrastructure ~ Code

TDD Cycle

Add Test

Refactor

Watch
Test Fail

Run tests

Write Code

Testing ~ Monitoring

Cucumber-nagios



Debugging ~ Metrics

Metrics Reuse

Traditional Ops

TEST


PROD




Collectd, Ganglia,
Graphite, Opentsdb

Extend “metrics” to Development


DEV



TEST




PROD




Collectd, Ganglia,
Graphite, Opentsdb

Extend “logs” to Development


DEV



TEST




PROD




Logstash, Graylog

Selfservicing “metrics injection”


DEV



TEST




PROD




StatsD

Selfservicing “alerts”


DEV



TEST




PROD




Tatle

Selfservicing “graphs”


DEV



TEST



PROD



Graphite

App
Deployment

Continuous
Integration

Load
Testing

App
Metrics

Monitoring

Alerting

App
Metrics

Logging

Heroku

Cloudbees

Blitz.io

New Relic

Pingdom

Pager Duty

DatadogHQ

Loggly



Repeating Service Pattern
Simple API/CLI , Self-Servicing

+Reuse workflow across monitoring tools

Collectd
Nagios

Ganglia
Zenoss

Graphite
Sensu

“my dream” Library

```
graph TD; Library["my dream Library"] --> Nagios; Library --> Zenoss; Library --> Sensu;
```

The diagram illustrates a central concept, "my dream Library", which is represented by a white text label at the bottom. Three white arrows originate from this central label and point upwards to three distinct monitoring tool stacks. The left stack consists of "Collectd" and "Nagios". The middle stack consists of "Ganglia" and "Zenoss". The right stack consists of "Graphite" and "Sensu". All text is in a bold, white, sans-serif font against a black background.

Workflow reuse

Monitoring UP

Abstracting

It's all events

Timestamp - Key - Value

Logs Metrics Monitoring Meta Ops

Social IT

radiate information to where it's needed

Business 'Pulse'



So maybe
you don't have
unlimited resources



But you can shape your
internal IT as a cloud



Questions?

Thank you!

