

Building GoDaddy.com's Compute Cloud

- Darren Shepherd, OSCON 2012

About Me – Darren Shepherd

- Long time Linux user (since 1998?)
- Absolutely love Xen and Virtualization
- Happen to program Java
- Working for GoDaddy.com for about 2 years (since 2010)
 - Lead back-end developer for Cloud Servers
- Prefers writing code over writing slides

Overview

- GoDaddy.com
- Cloud Servers Product
- Software Architecture
- Messaging/Orchestration

GoDaddy.com

- Go Daddy is the world's largest domain name registrar and Web hosting provider
 - More than 53 million domain names under management
- More than 10.4 million customers
- DNS, SSL, Web Hosting, VPS, Dedicated, Cloud Servers
- Go Daddy has more than 40 product offerings

Go Daddy Cloud Servers - Product

- Focus on usability and simplicity
- Target SMB's, small development shops, etc.

Go Daddy Cloud Servers - Details

- Servers

- Up to 16GB of RAM
- 40GB root disk
- Up to 1.2TB of additional storage can be added
- Ubuntu 12.04, Fedora 16, CentOS 5.8/6.2, Windows Server 2008R2

- Networking

- Private L2
- Assign multiple IP's to a network
- Load Balancing and Port Forwarding
- Source IP filtering
- VPN

- Storage

- Snapshot/Restore
- Volume from Snapshot

Control Panel

Network Assets

LIBRARY

- Storage
- Templates
- IP Groups
- Users

NETWORKS

- oneiric
 - boot
 - oneiric
 - openstack
 - tcbuild**
 - tctest
 - ubu
 - win
- prod
 - prodweb

tcbuild

Template: Ubuntu 12.04 - Base - 64-bit (version 20120611)

Attach Add

Elastic Virtual Disks

tcbuild Root [Create Backup](#)

Backups:

Jul 3, 2012

/dev/xvda

tc [Destroy](#) [Detach](#) [Create Backup](#)

Backups:

Jul 8, 2012 Today

/dev/xvdb 40 GB

Network Rules

[Add](#)

50.███

:2323 → :22 TCP SSH 1 Rule, Deny Others	ssh 	Edit Delete
:8081 TCP Unknown 1 Rule, Deny Others	web 	Edit Delete

Console

```
...  
tcbuild:~#
```

Machine Details

State:	Running
Auto Restart:	On
Private IP:	10.1.0.9
RAM:	1 GB
CPU:	4 vCPUs

[Edit Machine](#)

REST API

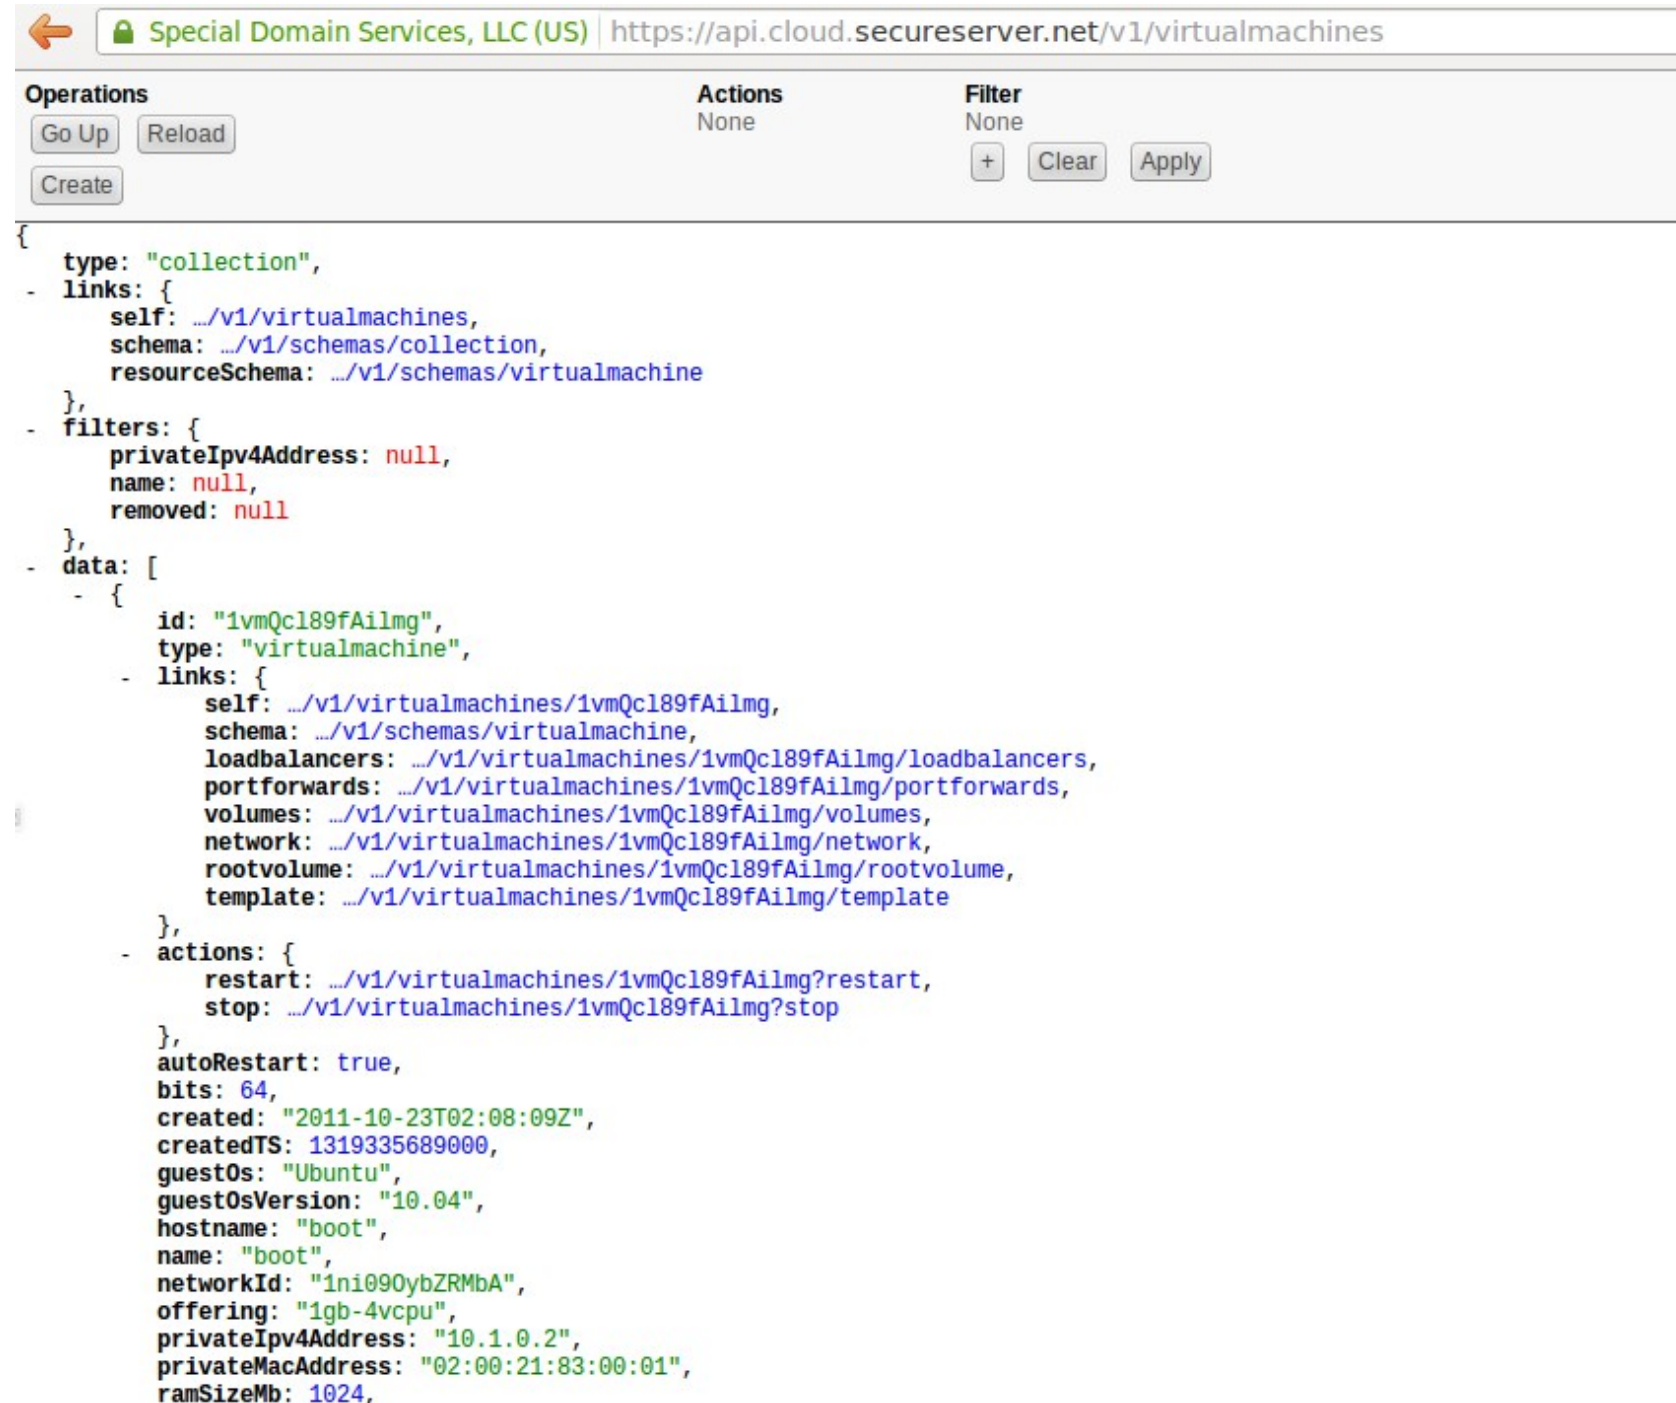
- “Curl Compatible”
- JSON/HTTPS
- HTTP Basic Auth
- Friendly URL's
 - /v1/virtualmachines
 - /v1/virtualmachines/42/volumes
- Hypermedia as the Engine of Application State (HATEOAS)
- Introspection capabilities through /v1/schemas
- HTML REST API UI
- URL encoded and multipart forms support

API - Raw Response

← Special Domain Services, LLC (US) https://api.cloud.secureserver.net/v1/virtualmachines?_format=json ☆ ▾

```
{
  "type" : "collection",
  "links" : {
    "self" : "https://api.cloud.secureserver.net/v1/virtualmachines?_format=json",
    "schema" : "https://api.cloud.secureserver.net/v1/schemas/collection",
    "resourceSchema" : "https://api.cloud.secureserver.net/v1/schemas/virtualmachine"
  },
  "filters" : {
    "privateIpv4Address" : null,
    "name" : null,
    "removed" : null
  },
  "data" : [ {
    "id" : "lvmQcl89fAilmg",
    "type" : "virtualmachine",
    "links" : {
      "self" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg",
      "schema" : "https://api.cloud.secureserver.net/v1/schemas/virtualmachine",
      "loadbalancers" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg/loadbalancers",
      "portforwards" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg/portforwards",
      "volumes" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg/volumes",
      "network" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg/network",
      "rootvolume" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg/rootvolume",
      "template" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg/template"
    },
    "actions" : {
      "restart" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg?restart",
      "stop" : "https://api.cloud.secureserver.net/v1/virtualmachines/lvmQcl89fAilmg?stop"
    },
    "autoRestart" : true,
    "bits" : 64,
    "created" : "2011-10-23T02:08:09Z",
    "createdTS" : 1319335689000,
    "guestOs" : "Ubuntu",
    "guestOsVersion" : "10.04",
    "hostname" : "boot",
    "name" : "boot",
    "networkId" : "lni090ybZRMbA",
    "offering" : "lgb-4vcpu",
    "privateIpv4Address" : "10.1.0.2",
    "privateMacAddress" : "02:00:21:83:00:01",
    "ramSizeMb" : 1024,
    "removed" : null,
    "removedTS" : null,
  } ]
}
```

API – When browser is detected



```
{
  type: "collection",
  - links: {
    self: .../v1/virtualmachines,
    schema: .../v1/schemas/collection,
    resourceSchema: .../v1/schemas/virtualmachine
  },
  - filters: {
    privateIpv4Address: null,
    name: null,
    removed: null
  },
  - data: [
    - {
      id: "1vmQc189fAilmg",
      type: "virtualmachine",
      - links: {
        self: .../v1/virtualmachines/1vmQc189fAilmg,
        schema: .../v1/schemas/virtualmachine,
        loadbalancers: .../v1/virtualmachines/1vmQc189fAilmg/loadbalancers,
        portforwards: .../v1/virtualmachines/1vmQc189fAilmg/portforwards,
        volumes: .../v1/virtualmachines/1vmQc189fAilmg/volumes,
        network: .../v1/virtualmachines/1vmQc189fAilmg/network,
        rootvolume: .../v1/virtualmachines/1vmQc189fAilmg/rootvolume,
        template: .../v1/virtualmachines/1vmQc189fAilmg/template
      },
      - actions: {
        restart: .../v1/virtualmachines/1vmQc189fAilmg?restart,
        stop: .../v1/virtualmachines/1vmQc189fAilmg?stop
      },
      autoRestart: true,
      bits: 64,
      created: "2011-10-23T02:08:09Z",
      createdTS: 1319335689000,
      guestOs: "Ubuntu",
      guestOsVersion: "10.04",
      hostname: "boot",
      name: "boot",
      networkId: "1ni090ybZRMbA",
      offering: "1gb-4vcpu",
      privateIpv4Address: "10.1.0.2",
      privateMacAddress: "02:00:21:83:00:01",
      ramSizeMb: 1024,
    }
  ]
}
```

What is this thing made of?

- Started development 2 years ago with Cloudstack
- Went live one year ago with a mostly Cloudstack cloud
 - UI and Storage changed
- Over the last year as we optimized and tailored Cloudstack we've taken a different architectural direction
- At this point most of Cloudstack has been replaced
- Future plans – either
 - Open source what we have
 - Converge with a another open source cloud stack

Software Stack

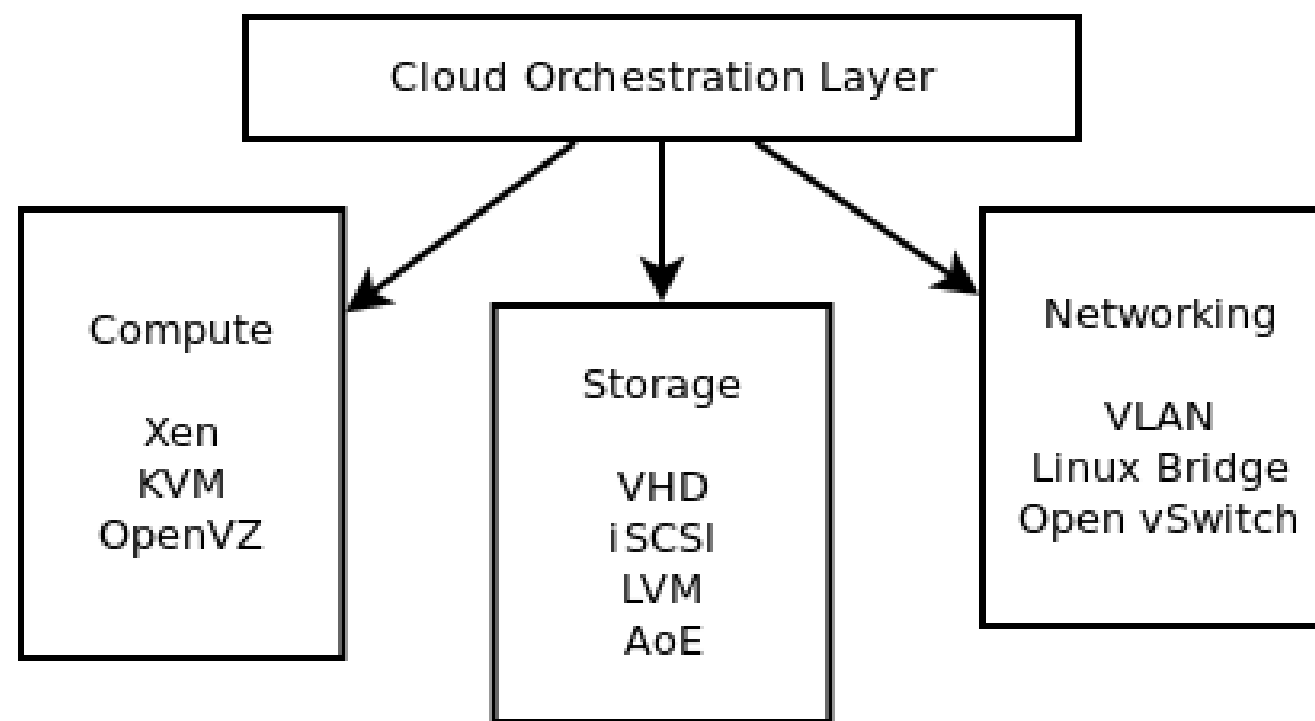
- PHP – UI
- Java – Backend
 - Spring, Hibernate, Jackson, and tons of Apache libraries
- Node.js – Network oriented services or agents

System Engineering

- Compute
 - Citrix XenServer – PVGRUB and HVM for Windows
- Networking
 - VLAN's
 - Custom VLAN bridging appliance
- Storage
 - VHD on NFS

Basics of a Cloud Compute Provisioning Platform

- Other technologies provide implementation of Compute, Storage, Networking
- A good platform should at least
 - Provide a good abstraction model
 - Reliably provision resources
 - Recover when possible
 - Fail gracefully when possible
 - Or just fail otherwise

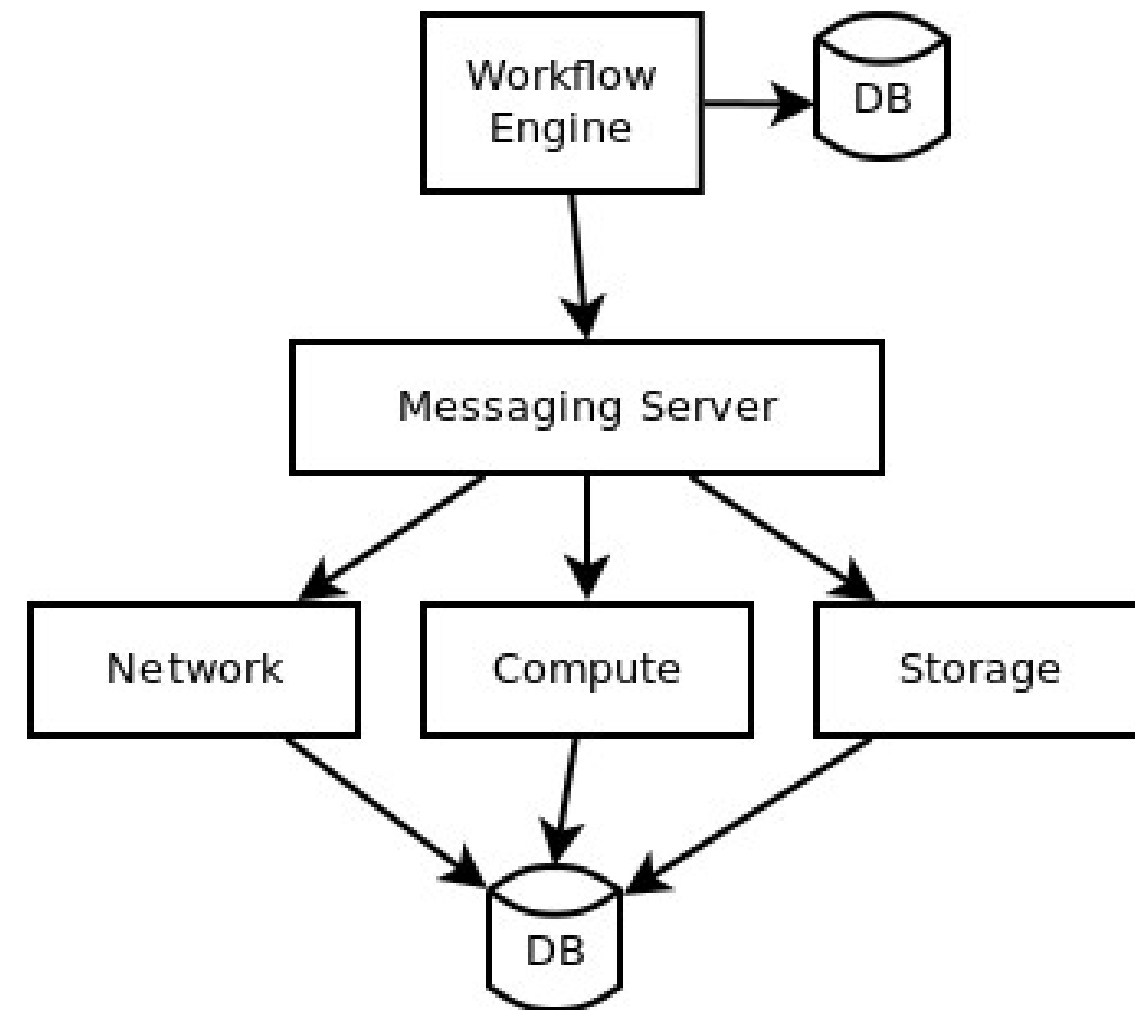


We had two approaches for service orchestration

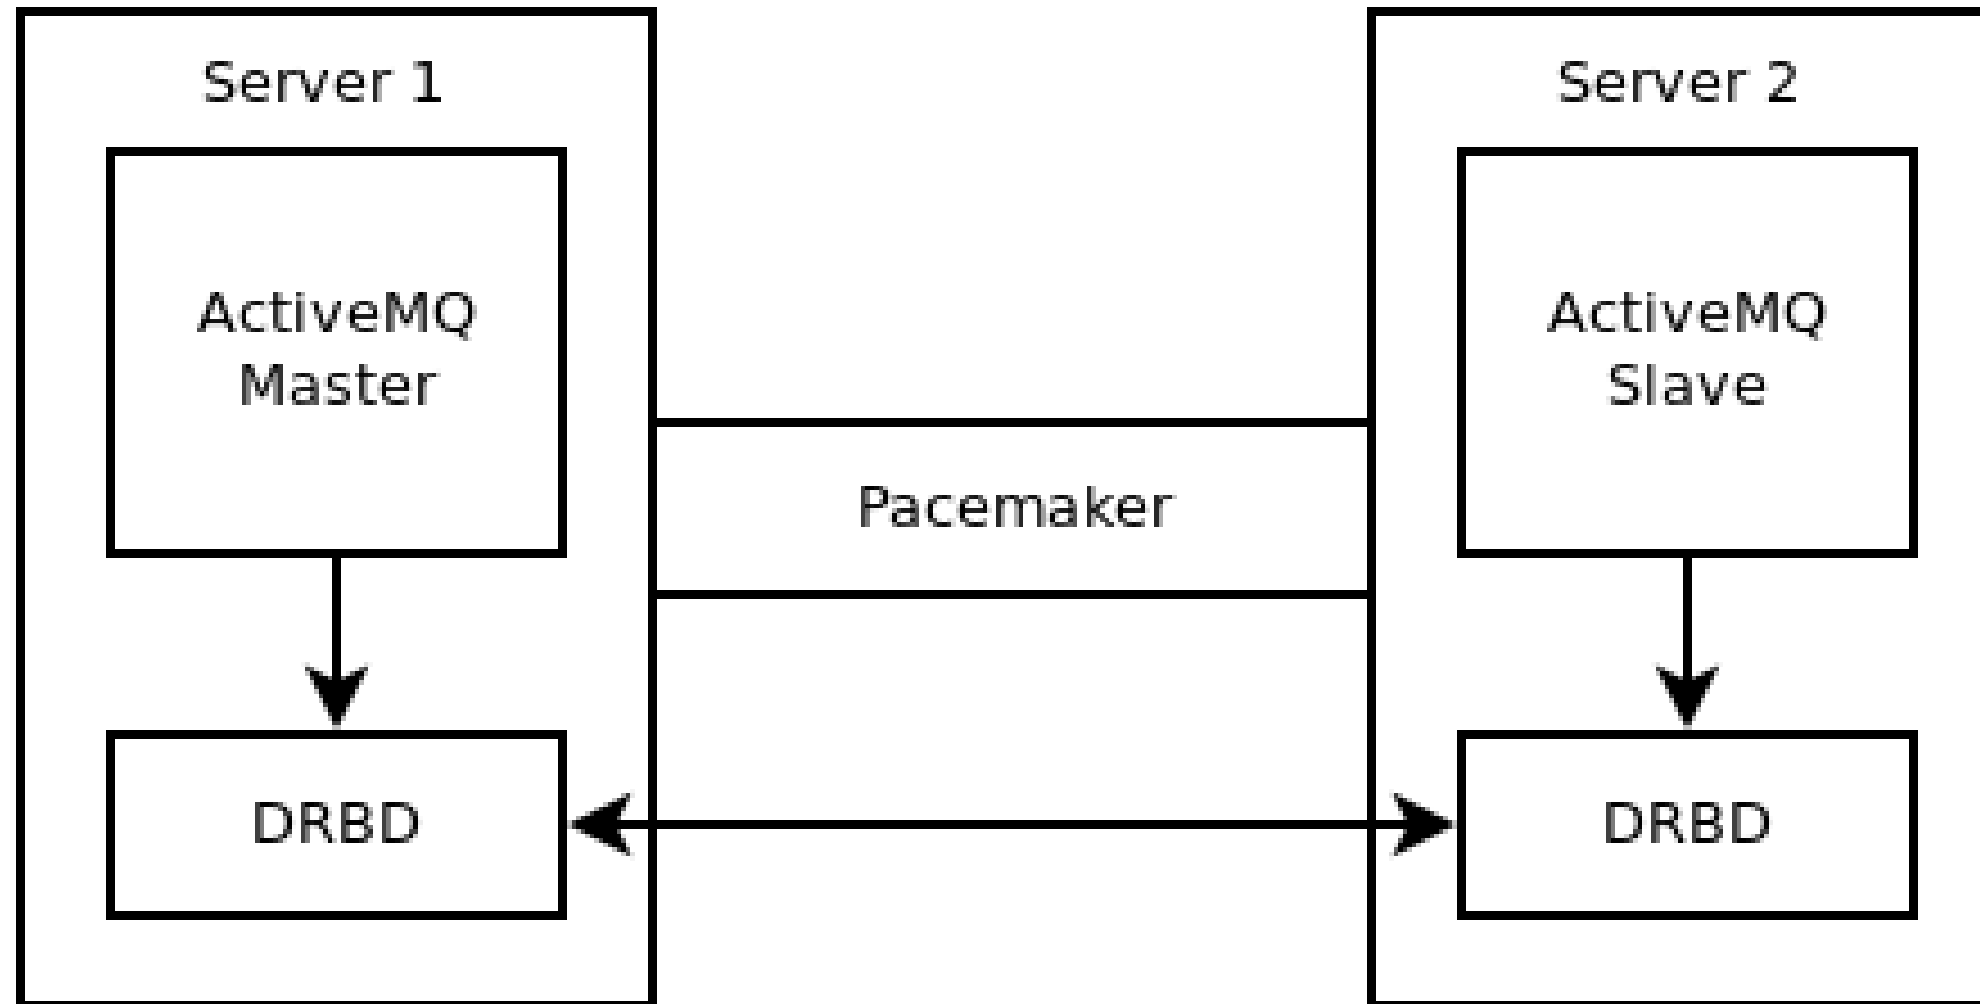
Synchronous

```
create_volume();  
setup_network();  
start_vm();
```

JMS Based Workflow



Our Apache ActiveMQ HA Setup



Our Event Driven Architecture

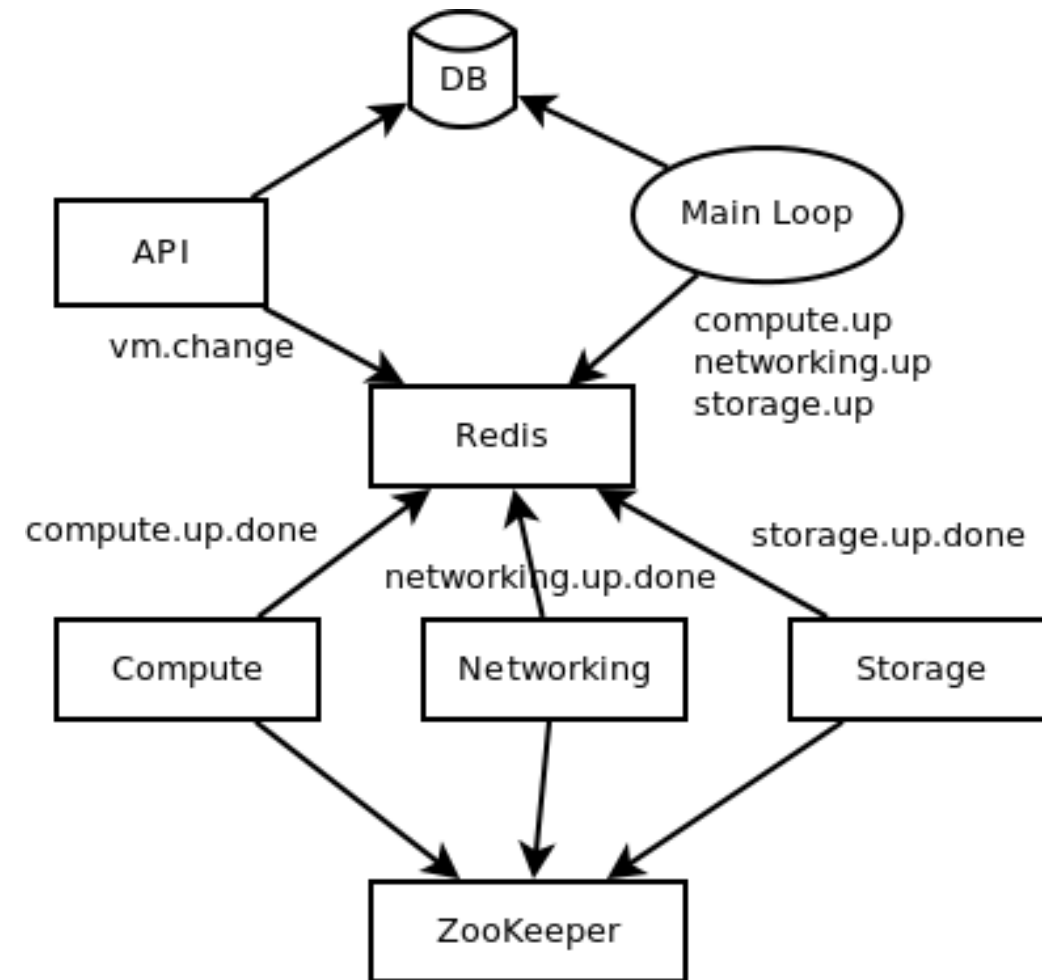
- No guarantee that events will be received
 - Events should be idempotent
 - Events should not be queued
 - Events should be resent until state is reconciled
- System is series of states that must be reconciled
 - Desired state should be recorded
 - Services update state or send events indication state should be updated
- Locks ensure no event single is processed concurrently
- Crash-only design

Redis & Zookeeper

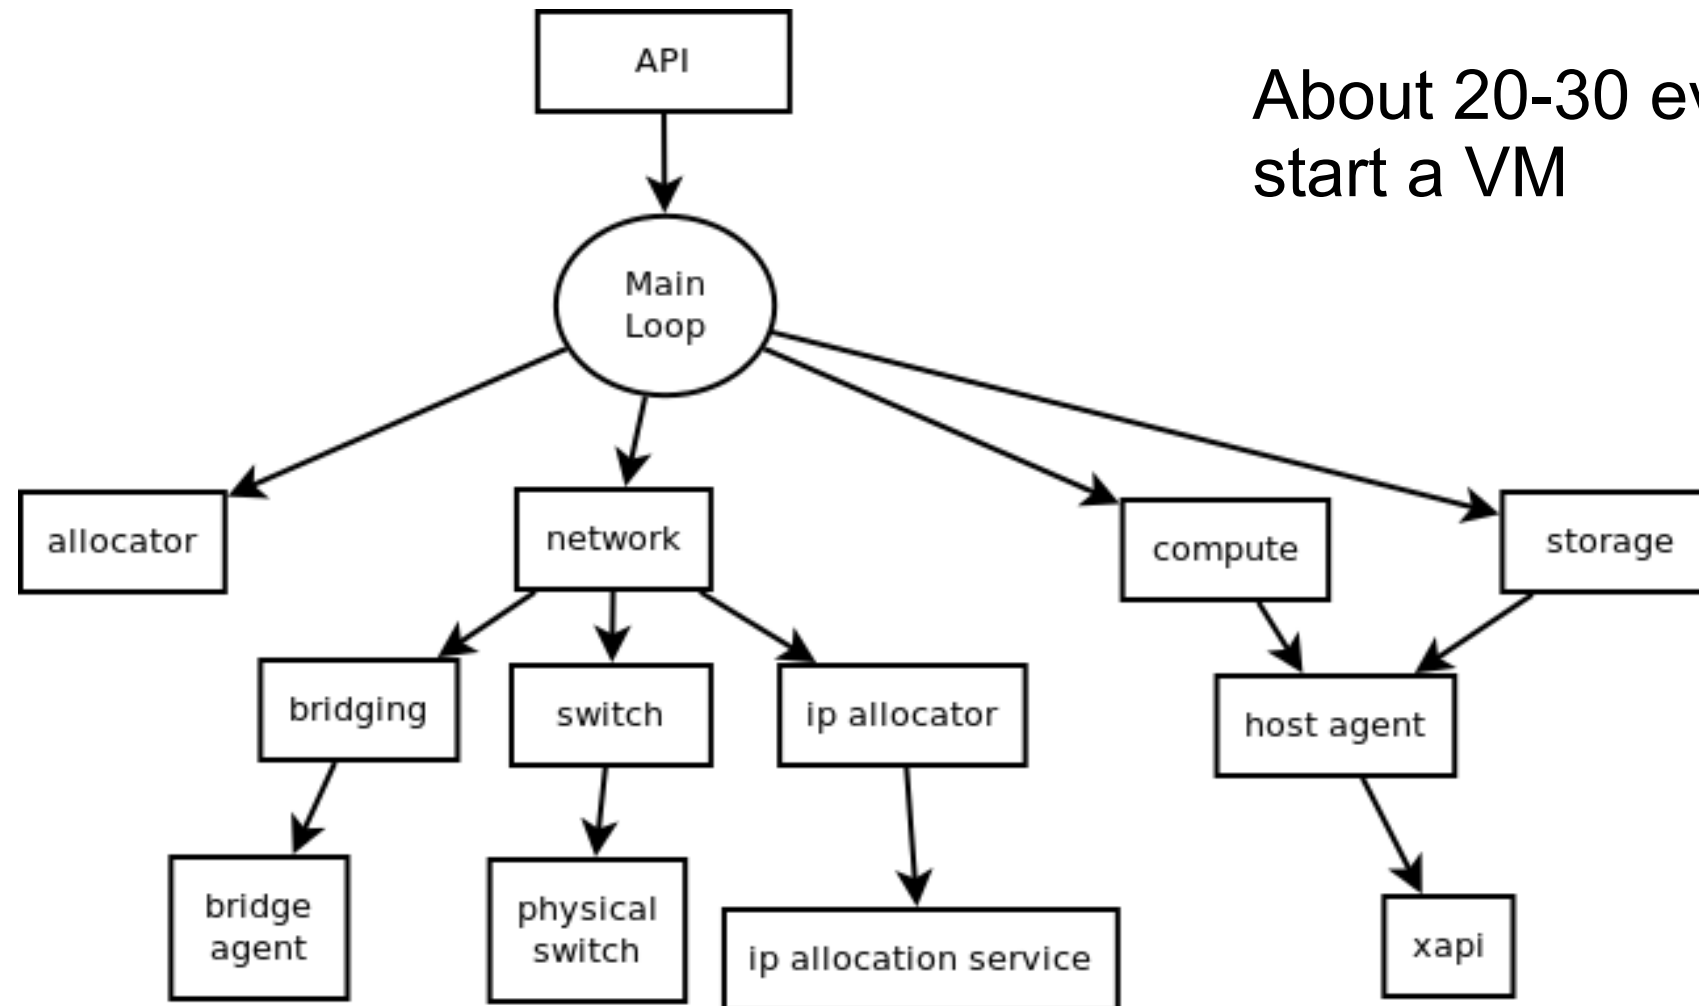
- Redis provides basic PUB/SUB
 - Provides almost no features, which is great
- Zookeeper works as a distributed lock manager

Main loop

- API service records requested state to DB
 - Sends Redis event to process change
- Main Loop reads state from DB on interval or triggered from event
- Main Loop sends events to other systems



Real Architecture



About 20-30 events are fired to start a VM

Done, let's have lunch... or Q&A