
BARCELONA
DRUPALCON 2015

Migrating a running service to AWS

Nick Veenhof

Ricardo Amaro

DevOps Track

https://events.drupal.org/barcelona2015/sessions/migrating-running-service-mollom-aws-without-service-interruptions-and-reduce

The Developer

@Nick_vh

Ghent

Barcelona

Lisbon

Boston

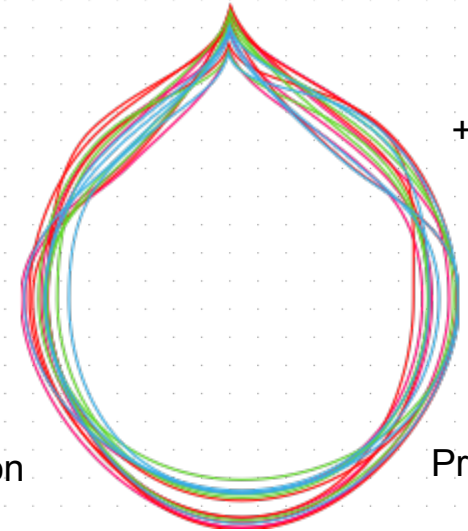

+8 Years in Drupal

Search++

4 years at Acquia

Principal Software Engineer

So good to be back...

Mollom

- Detecting Spam from Ham
  - Reducing your moderation efforts
- Very fast response times (avg under 50 msec)
- Fully Managed SAAS service
- Free and paid version
- Downtime means unprotected sites, which is bad for reputation and adoption
- Built in Java
-

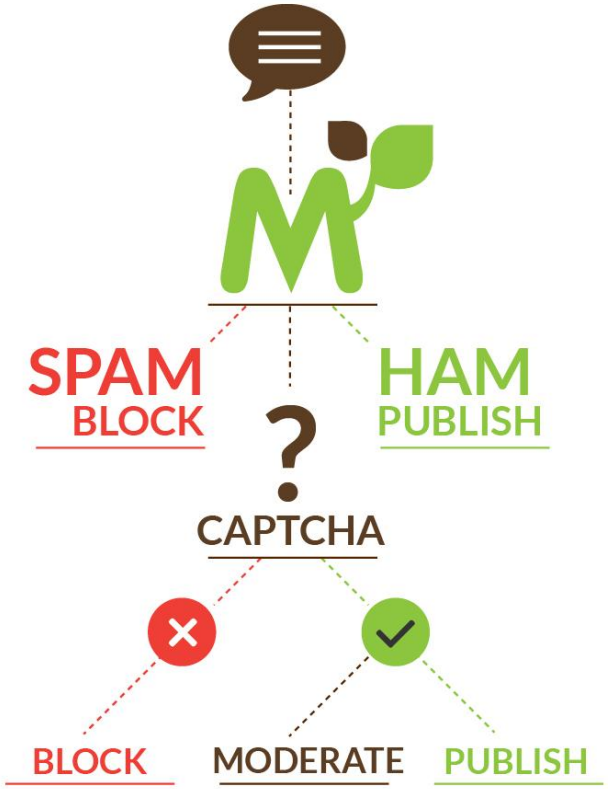

The Opsian

@ricardoamaro

Portugal

Lisbon

Family

Drupal Community

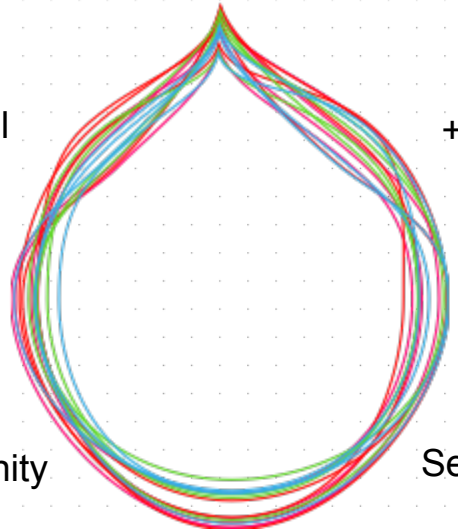

+7 years Drupal

90’s Linux Adopter

4 years at Acquia

Senior Tier2 Ops Engineer

Pre-Migration

Roses, Roses everywhere...

How we got the news…

*”Operations is now responsible for Mollom servers being **up or down**, and basic services being available (such as SSH, apache, nginx, etc). If further problems persist above the services layer into the application layer, Ops is to escalate to Mollom Engineering immediately. “*

Highly complex piece of engineering on top of non-cloud hosting.

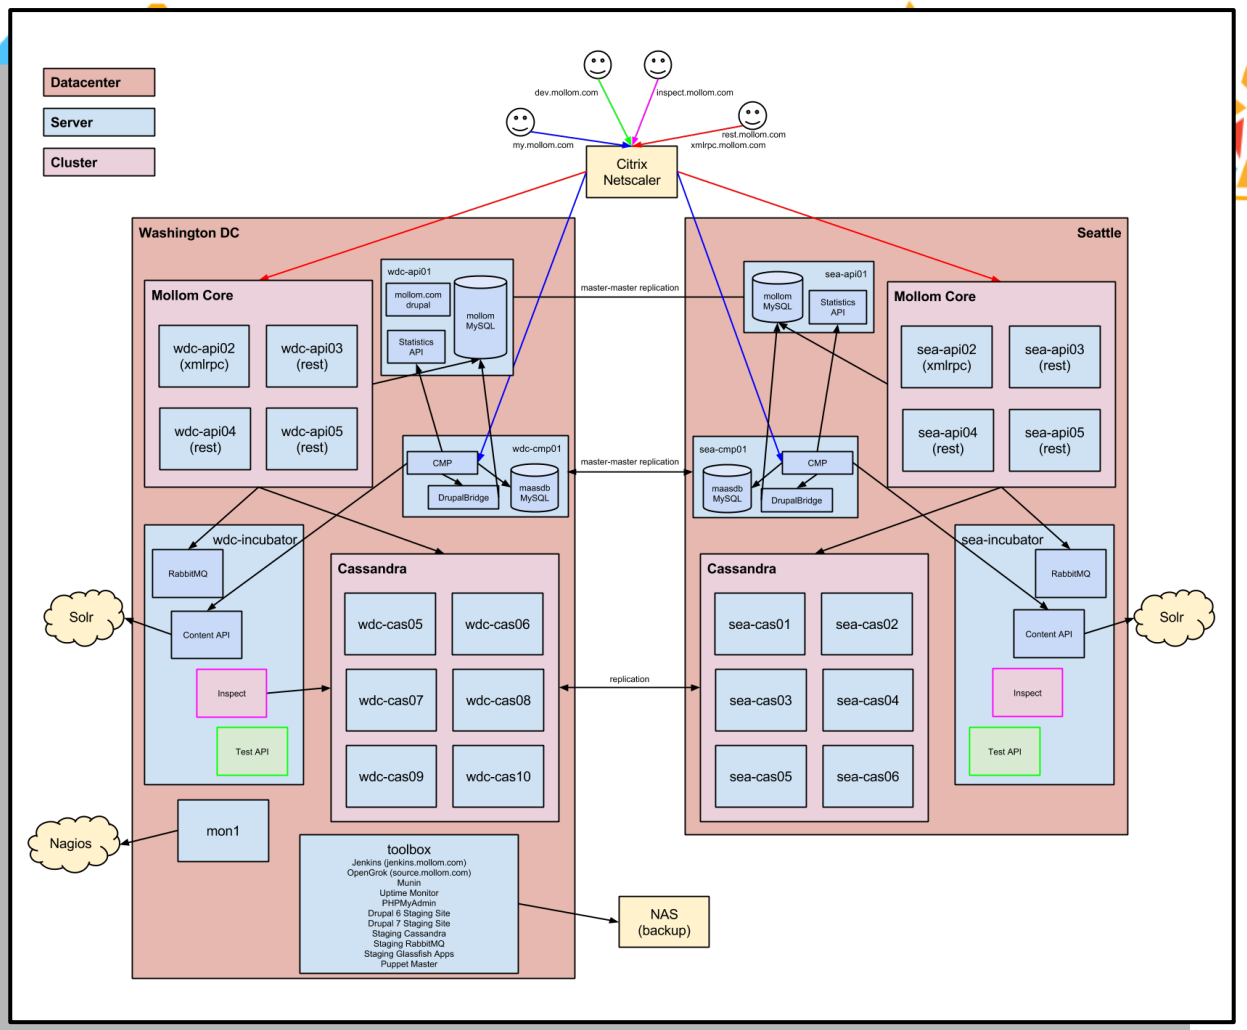

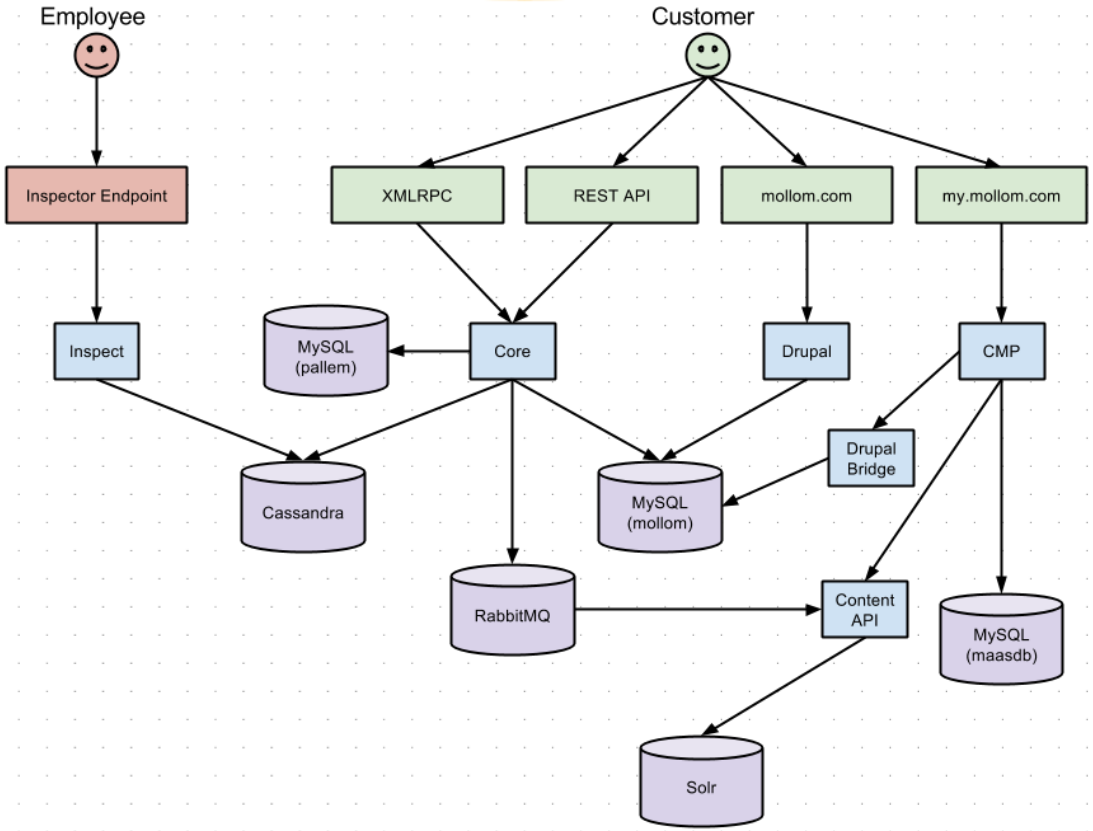
Employee
Customer
Inspector Endpoint
XMLRPC
REST API
mollom.com
my.mollom.com
Inspect
MySQL
(pallem)
Core
Drupal
CMP
Drupal
Bridge
Cassandra
MySQL
(mollom)
RabbitMQ
Content
API
MySQL
(maasdb)
Solr


DON'T WORRY HUMAN
I FIX PROBLEM

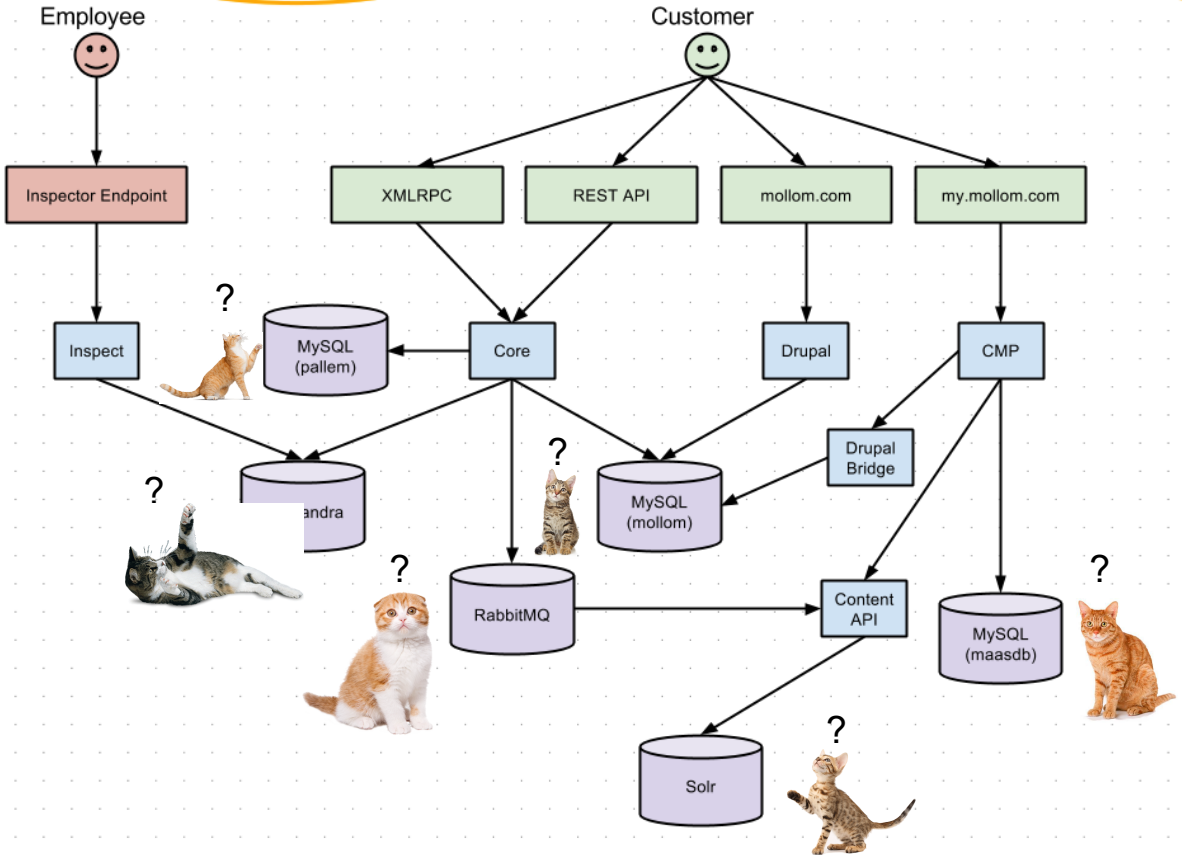
Employee
Customer
Inspector Endpoint
XMLRPC
REST API
mollom.com
my.mollom.com
Inspect
MySQL (pallem)
Core
Drupal
CMP
Drupal Bridge
MySQL (mollom)
RabbitMQ
Content API
MySQL (maasdb)
Solr

20 million http requests per day

8 million of spam requests / day

worst day: 300+ alerts...

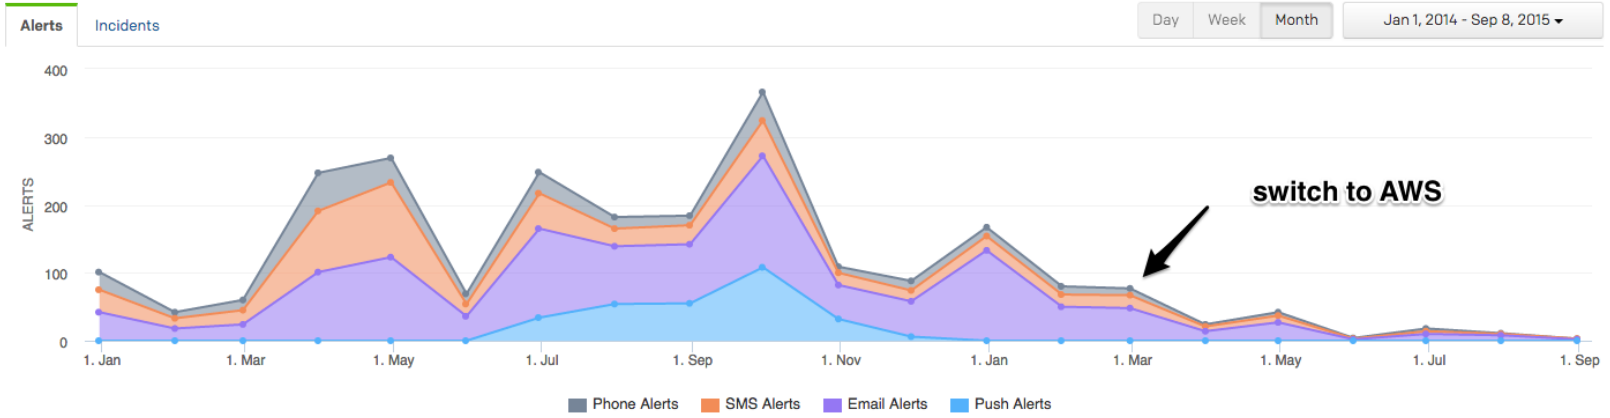

One clear guidance example...

Question: “Is disk usage above 95%?”

Answer: “Remove all files that start with the same prefix as the data file...”

```
rm -rf Mollom-session_history-he-78609-*
```

“... and restart Cassandra”

```
/etc/init.d/cassandra restart
```

Architecture Exercise

Look before you leap

Exercise

- One row = One Component.
- I need to be able to “take down” someone and still be up and running
- Order is important. I will be a site visitor, so I want you to start from the front to the end.

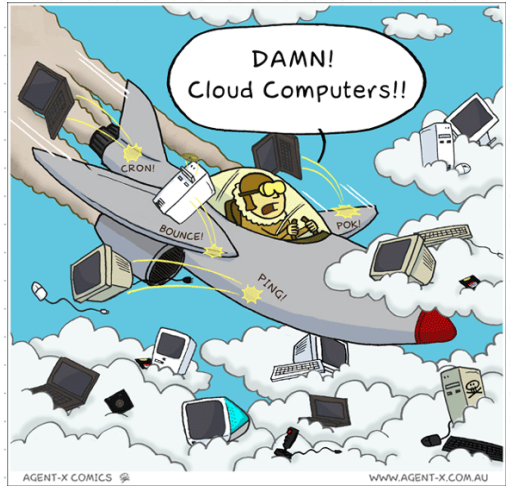

Exercise

- Reverse Proxy (VARNISH)
- Web Server (WEB)
- DNS
- Load Balancer (LB)
- Database (DB)
- Object Caching (Cache)


Ephemeralism

Eye-opener

The Practice of Cloud System Administration

Describes the optimal environment and how this relates to reality. Warning, there is no perfect.

A very digestible book for designing distributed systems. This book exposes software patterns that every cloud infrastructure engineer should know.

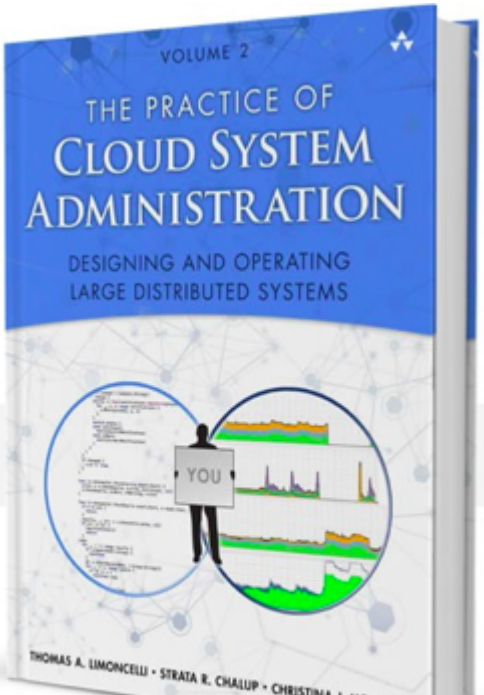

CAP Theorem

The Practice of Cloud System Administration

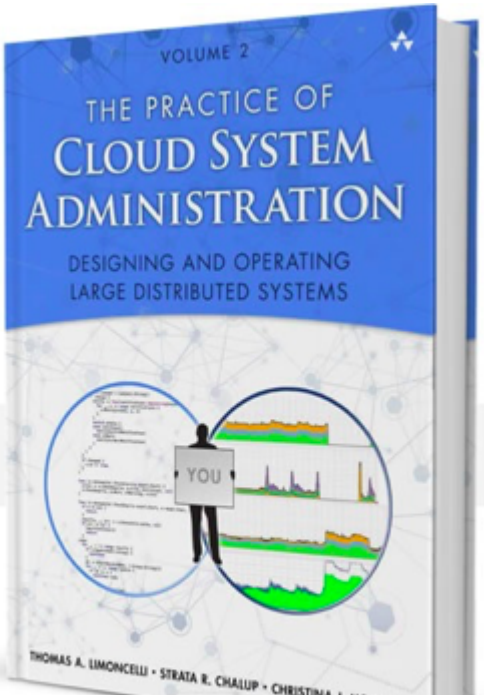

It is impossible for a distributed computer system to simultaneously provide all three of the following guarantees:

- *Consistency* (all nodes see the same data at the same time)
- *Availability* (a guarantee that every request receives a response about whether it succeeded or failed)
- *Partition tolerance* (the system continues to operate despite arbitrary partitioning due to network failures)

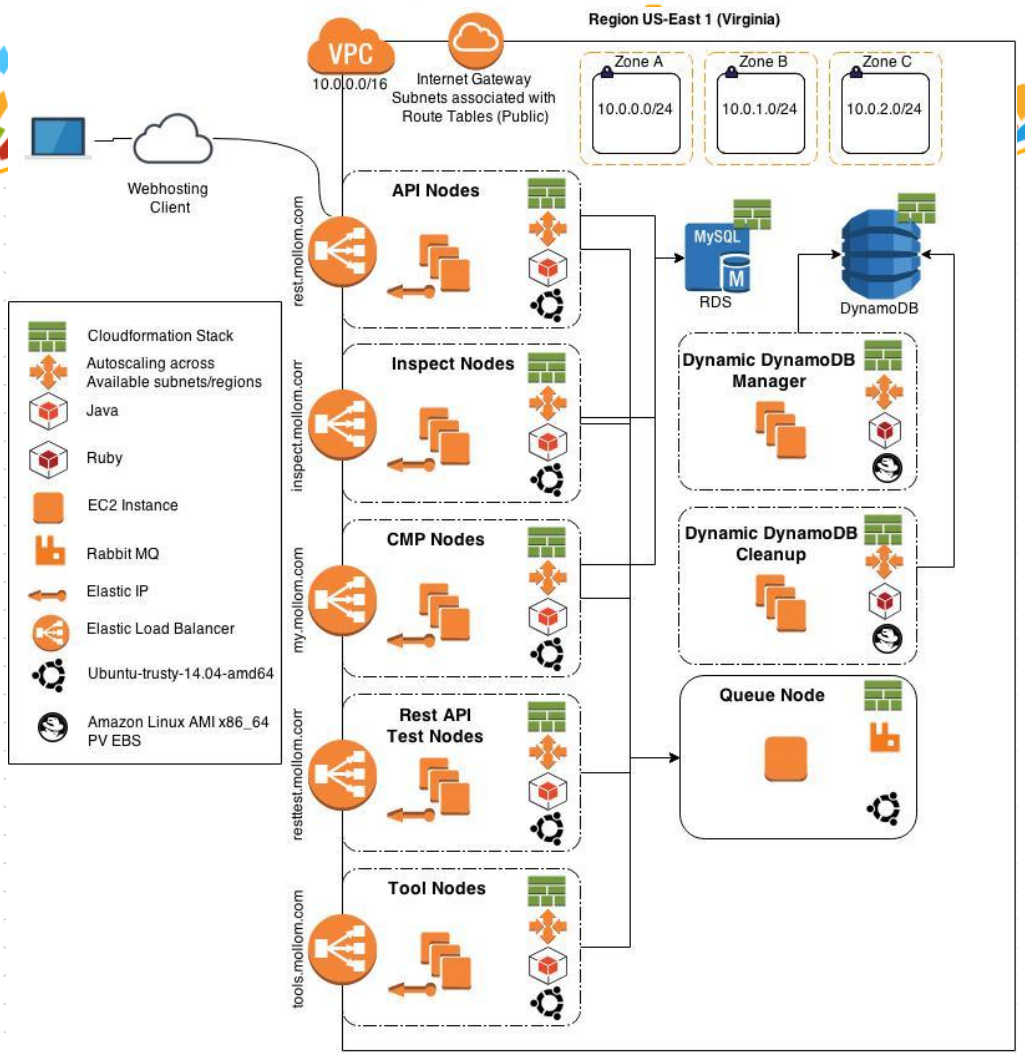
Region US-East 1 (Virginia)
VPC
10.0.0.0/16
Internet Gateway
Subnets associated with
Route Tables (Public)
Zone A
10.0.0.0/24
Zone B
10.0.1.0/24
Zone C
10.0.2.0/24
Webhosting
Client
API Nodes
rest.mollom.com
Inspect Nodes
inspect.mollom.com
CMP Nodes
my.mollom.com
Rest API
Test Nodes
resttest.mollom.com
Tool Nodes
tools.mollom.com
MySQL
RDS
DynamoDB
Dynamic DynamoDB
Manager
Dynamic DynamoDB
Cleanup
Queue Node
Cloudformation Stack
Autoscaling across
Available subnets/regions
Java
Ruby
EC2 Instance
Rabbit MQ
Elastic IP
Elastic Load Balancer
Ubuntu-trusty-14.04-amd64
Amazon Linux AMI x86_64
PV EBS


WELL, ACTUALLY

Cloudformation

Stackin' it up

"AWS **CloudFormation** is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS."

AWS
CloudFormation

Cloudformation

Stackin' it up

- AutoScaling Groups (ASG)
- Elastic Load Balancer (ELB)
- Elastic Compute 2 (EC2)
- AMI (VM of Ubuntu 14.04)
- Java

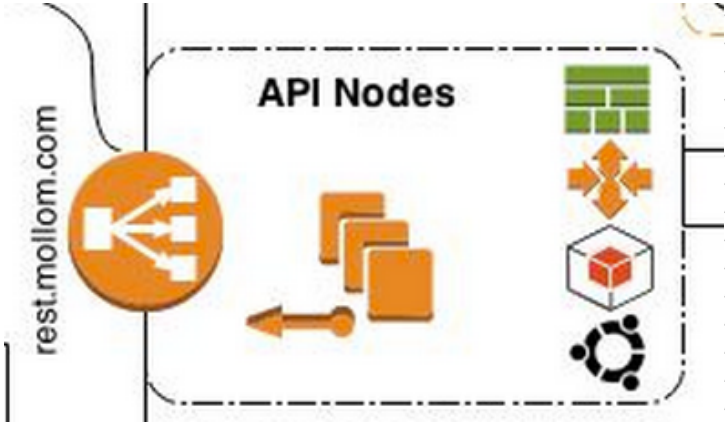

Cloudformation

```
"LoadBalancer": {
  "Type": "AWS::ElasticLoadBalancing::LoadBalancer",
  "Properties": {
    "AccessLoggingPolicy": {
      "EmitInterval": 60,
      "Enabled": true,
      "S3BucketName": {
        "Fn::Join": [ ".",
          [
            "[redacted]",
            { "Ref": "AWS::Region" },
            { "Ref": "EnvironmentName" },
            "logs"
          ]
        ]
      },
      "S3BucketPrefix": "[redacted]"
    },
    "HealthCheck": {
      "HealthyThreshold" : "4",
      "Interval" : "30",
      "Target" : "HTTP:8080/ping",
      "Timeout" : "5",
      "UnhealthyThreshold" : "2"
    },
    "CrossZone" : "true",
    "ConnectionDrainingPolicy": {
      "Enabled" : "true",
      "Timeout" : "10"
    },
    "SecurityGroups": [ { "Ref" : "ELBSecurityGroup" } ],
```

```
{
  "AWSTemplateFormatVersion": "2010-09-09",
  "Description": "AWS CloudFormation template for setting up Dice Proxy api.",
  "Parameters": {
    "Vpc" : {
      "Type" : "String",
      "Description" : "Id of VPC"
    },
    "AcquiaUbuntuMirrorVersion": {
      "Type": "String",
      "Description": "Version of the mirror we want to use. See [redacted] for a list of all the versions.",
      "Default": "2015-07-13"
    },
    "Route53Domain" : {
      "Type" : "String",
      "Description" : "Domain to which we can register our domain to."
    },
    "PublicSubnets" : {
      "Type" : "CommaDelimitedList",
      "Description" : "List of VPC Subnets where the stack will be launched"
    },
    "[redacted]" : {
      "Type" : "String",
      "Description" : "[redacted]"
    },
```

Virtual Private Cloud (VPC)

Isolation isn't bad, mkay?

Amazon VPC lets you provision a logically isolated section of the Amazon Web Services (AWS) Cloud where you can launch AWS resources in a virtual network that you define.

Virtual Private Cloud (VPC)

Isolation isn't bad, mkay?

- Private Subnets
- Internal Load Balancers
- Public IP addresses
- Security Groups

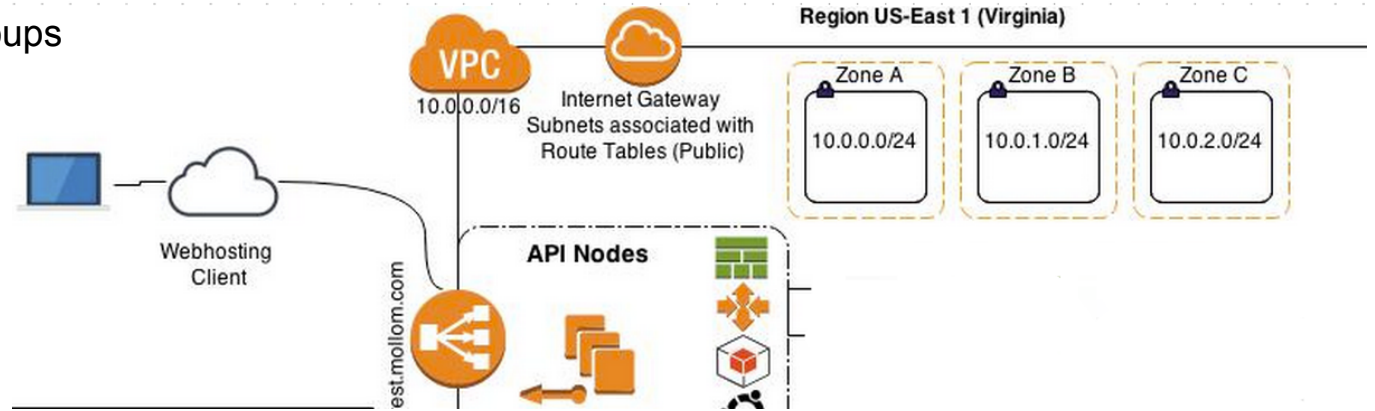

Virtual Private Cloud (VPC)

Isolation isn't bad, mkay?

```
"InternetGateway" : {
  "Type" : "AWS::EC2::InternetGateway",
  "Properties" : {
    "Tags" : [
      { "Key" : "Application", "Value" : { "Ref" : "AWS::StackId" } },
      { "Key" : "Network", "Value" : "Public" }
    ]
  }
},

"GatewayToInternet" : {
   "Type" : "AWS::EC2::VPCGatewayAttachment",
   "Properties" : {
     "VpcId" : { "Ref" : "Vpc" },
     "InternetGatewayId" : { "Ref" : "InternetGateway" }
   }
},

"PublicRouteTable" : {
  "Type" : "AWS::EC2::RouteTable",
  "Properties" : {
    "VpcId" : { "Ref" : "Vpc" },
    "Tags" : [
      { "Key" : "Application", "Value" : { "Ref" : "AWS::StackId" } },
      { "Key" : "Network", "Value" : "Public" }
    ]
  }
},
```

```
"PublicRoute" : {
  "Type" : "AWS::EC2::Route",
  "DependsOn" : "GatewayToInternet",
  "Properties" : {
    "RouteTableId" : { "Ref" : "PublicRouteTable" },
    "DestinationCidrBlock" : "0.0.0.0/0",
    "GatewayId" : { "Ref" : "InternetGateway" }
  }
},

"PublicSubnetRouteTableAssociation1" : {
  "Type" : "AWS::EC2::SubnetRouteTableAssociation",
  "Properties" : {
    "SubnetId" : { "Ref" : "PublicSubnet1" },
    "RouteTableId" : { "Ref" : "PublicRouteTable" }
  }
},

"PublicSubnetRouteTableAssociation2" : {
  "Type" : "AWS::EC2::SubnetRouteTableAssociation",
  "Properties" : {
    "SubnetId" : { "Ref" : "PublicSubnet2" },
    "RouteTableId" : { "Ref" : "PublicRouteTable" }
  }
},
```

Relational Database Service

It's not a triptych

- Fully Managed
- H/A possible
- Within your VPC, non public
- Option to use MariaDB, Postgres, Aurora, …
- Highly configurable

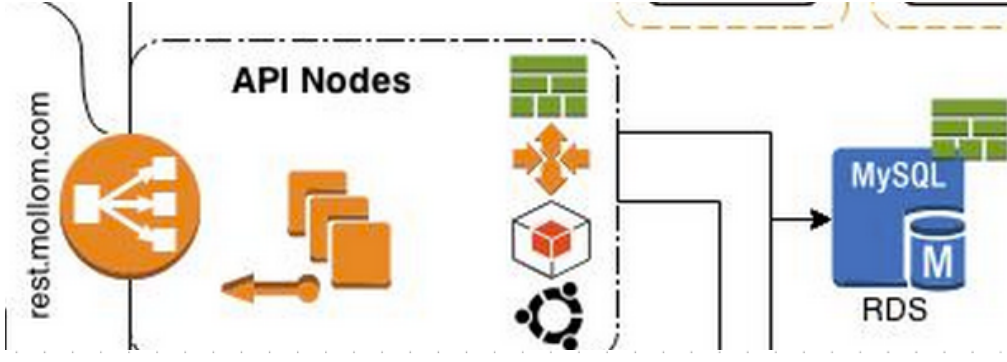

Relational Database Service

It's not a triptych

```
"ProductionDatabase": {
  "Type": "AWS::RDS::DBInstance",
  "Condition" : "IsProduction",
  "Properties": {
    "PubliclyAccessible" : "false",
    "VPCSecurityGroups": [
      {
        "Ref": "DatabaseSecurityGroup"
      }
    ],
    "DBSubnetGroupName": {
      "Ref": "dbSubnetGroup"
    },
    "DBInstanceClass": {
      "Ref": "DBInstanceType"
    },
    "DBParameterGroupName": {
      "Ref": "diceParams"
    },
    "AllocatedStorage": {
      "Ref": "DBStorage"
    },
    "Engine": "MySQL",
    "MasterUsername": {
      "Ref": "DBUsername"
    },
```

```
"DatabaseCPUAlarmWarning": {
  "Type" : "AWS::CloudWatch::Alarm",
  "Properties" : {
    "ActionsEnabled" : "False",
    "AlarmDescription" : "Checks the RDS CPU Utilization",
    "AlarmName" : {
      "Fn::Join": [
        "",
        [
          { "Ref" : "EnvironmentName" },
          ".rds.cpu_idle.warning"
        ]
      ]
    },
    "ComparisonOperator" : "GreaterThanThreshold",
    "Dimensions" : [
      {
        "Name" : "DBInstanceIdentifier",
        "Value" : {
          "Fn::If" : [
            "IsProduction",
            { "Ref": "ProductionDatabase" },
            { "Ref": "NonProductionDatabase" }
          ]
        }
      }
    ],
    "EvaluationPeriods" : "2",
    "MetricName" : "CPUUtilization",
    "Namespace" : "AWS/RDS",
    "Period" : "900",
    "Statistic" : "Average",
    "Threshold" : "80"
  }
```

DynamoDB

Datawarehousing for the masses

AWS says: “DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability.”

We read: Cassandra without maintenance (and serious reduction in alerts)!

DynamoDB

Document storage for the masses

- Really fast
- Fully Managed
- No TTL, so we use rotation based tables
- Pricy, but maintenance-free.

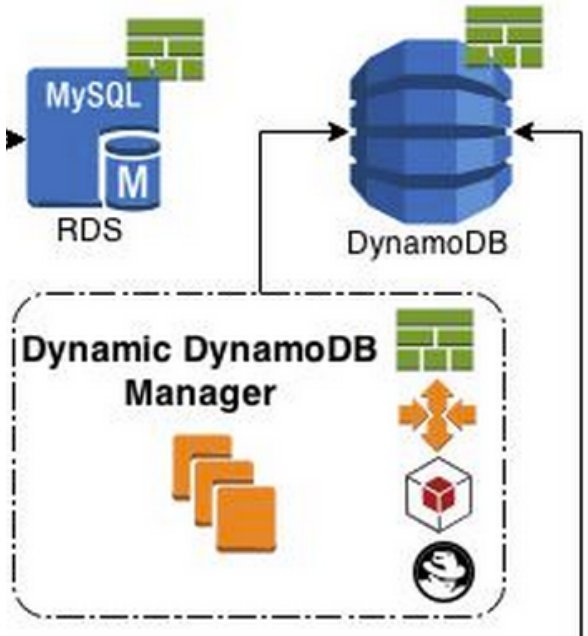

DynamoDB

Datawarehousing for the masses

- Dynamic DynamoDB
  - https://github.com/sebdah/dynamic-dynamodb
- Dynamic DynamoDB Manager
  - https://github.com/Mollom/dynamic-dynamodb-manager

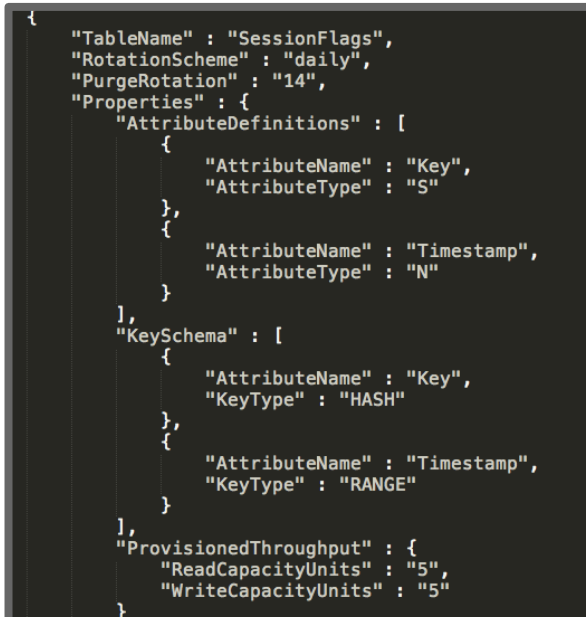

```
{
    "TableName" : "SessionFlags",
    "RotationScheme" : "daily",
    "PurgeRotation" : "14",
    "Properties" : {
        "AttributeDefinitions" : [
            {
                "AttributeName" : "Key",
                "AttributeType" : "S"
            },
            {
                "AttributeName" : "Timestamp",
                "AttributeType" : "N"
            }
        ],
        "KeySchema" : [
            {
                "AttributeName" : "Key",
                "KeyType" : "HASH"
            },
            {
                "AttributeName" : "Timestamp",
                "KeyType" : "RANGE"
            }
        ],
        "ProvisionedThroughput" : {
            "ReadCapacityUnits" : "5",
            "WriteCapacityUnits" : "5"
        }
```

EC2 + Load Balancing

VMception

Elastic Load Balancing (Amazon ELB) automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud.

EC2 = a VM, hosted on AWS's supervisor system.

Elastic Load Ba
distributes incor
Amazon EC2 in

EC2 = a VM, hc

EC2 + ELB

VMception

- Linux as you know it
- AMI-based
- Can disappear or crash. Don’t try to do non-stateless apps.
- Triggers to auto-scale (read: add/remove a ec2 machine) on predefined inputs.
- Update scheme involves disposable EC2 instances

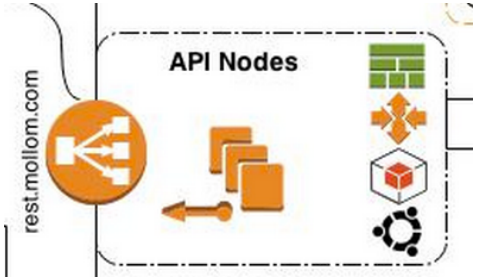

EC2 + ELB

Vmception

```
"AutoscalingGroup": {
  "Type": "AWS::AutoScaling::AutoScalingGroup",
  "UpdatePolicy": {
    "AutoScalingRollingUpdate": {
      "MinInstancesInService": "1",
      "MaxBatchSize": "1",
      "PauseTime" : { "Ref" : "PauseTime" }
    }
  },
  "Properties": {
    "AvailabilityZones": [
      { "Ref" : "AvailabilityZone1" },
      { "Ref" : "AvailabilityZone2" },
      { "Ref" : "AvailabilityZone3" }
    ],
    "VPCZoneIdentifier" : { "Ref" : "PublicSubnets" },
    "MaxSize": "100",
    "MinSize": "2",
    "DesiredCapacity": { "Ref" : "DesiredApiCapacity" },
    "LoadBalancerNames": [
      {
        "Ref": "LoadBalancer"
      }
    ],
    "LaunchConfigurationName": {
      "Ref": "LaunchConfiguration"
    },
    "Tags" : [ {
      "Key" : "Name",
      "Value" : "api-node",
      "PropagateAtLaunch" : "true"
    },
    {
      "Key" : "Environment",
```

```
"LaunchConfiguration": {
  "Type": "AWS::AutoScaling::LaunchConfiguration",
  "Metadata": {
    "Puppet" : {
      "roles" : [ "dice-api-node" ]
    },
    "AWS::CloudFormation::Init": {
      "config": {
        "packages": {
          "apt": {
            "ntp": [],
            "puppet": [],
            "mysql-client": [],
            "python-dev": [],
            "python-support": [],
            "syslog-ng-core": [],
            "syslog-ng": []
          },
          "python": {
            "boto": ["2.38.0"],
            "psutil": ["3.1.1"],
            "awscli": ["1.8.1"]
          }
        },
        "files": {
          "/opt/diamond/diamond.deb": {
            "source": {
              "Fn::Join": [
                "",
                [
                  "https://s3.amazonaws.com/",
                  {
                    "Ref": "ResourceBucket"
```

EC2 + ELB

Vmception

- Access Logging
- Health Check
- H/A (multiple zones)
- Connection Draining
- IPTables-like functionality
- Multiple listeners (read: port forwarding)
- SSL Termination (port 443, check cert and forward to HTTP port 80, eg SSL termination at the load balancer level)

```
"LoadBalancer": {
  "Type": "AWS::ElasticLoadBalancing::LoadBalancer",
  "Properties": {
    "AccessLoggingPolicy": {
      "EmitInterval": 60,
      "Enabled": true,
      "S3BucketName": {
        "Fn::Join": [ ".",
          [
            [redacted]
            { "Ref": "AWS::Region" },
            { "Ref": "EnvironmentName" },
            "logs"
          ]
        ]
      },
      "S3BucketPrefix": "api-proxy"
    },
    "HealthCheck": {
      "HealthyThreshold" : "4",
      "Interval" : "30",
      "Target" : "HTTP:8080/ping",
      "Timeout" : "5",
      "UnhealthyThreshold" : "2"
    },
    "CrossZone" : "true",
    "ConnectionDrainingPolicy": {
       "Enabled" : "true",
       "Timeout" : "10"
    },
    "SecurityGroups": [ { "Ref" : "ELBSecurityGroup" } ],
    "Subnets": {
      "Ref": "PublicSubnets"
    },
    "Listeners": [
      {
```

EC2 + ELB

So puppet or chef right?

- No puppet
- No Chef
- No Ansible
- Everything is fully rebuilt on launch, every update is a new machine
- We do not update single packages, we remove and add machines.
- Allows for returning to a point in time as the full “state” is preserved. Note: Data backups are still necessary if this is required.

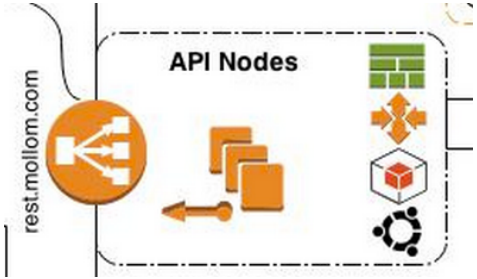

Metrics

Ever seen a cloud with a watch?

- AWS Cloudwatch
- Diamond + Custom Handlers
  - https://github.com/python-diamond/Diamond
- StatsD / Graphite
- Creating AWS Cloudwatch alarms
  per instance for non AWS-specific services

```
[[[GoMemoryGcTotalPause]]]
collector = gosystem
metric = MemoryGcTotalPause
name = go_memory_gctotalpause
unit = Seconds

[[[GoMemoryHeap]]]
collector = gosystem
metric = GoMemoryHeap
name = go_memory_heap
unit = Bytes

[[[GoMemoryBytesInStack]]]
collector = gosystem
metric = MemoryBytesInStack
name = go_memory_bytesinstack
unit = Bytes
```

```
[[[DiskspaceBytePercentFree]]]
collector = diskspace
metric = cloudimg-rootfs.byte_percentfree
name = disk_space_percentfree
unit = Percent
warning = 10
critical = 5
comparison = <
description = disk space percent free
evaluation_periods = 1
```

Alarms

Every Pager has its duty

- Nagios + Pagerduty
- Integration with Cloudwatch
- Ordering of alerts, to help those who are on-call to prioritize.

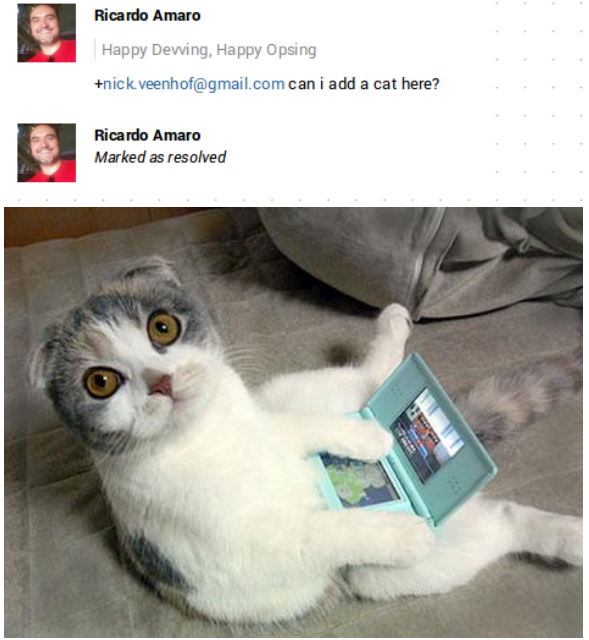

DNS

Returning a different IP based on your region

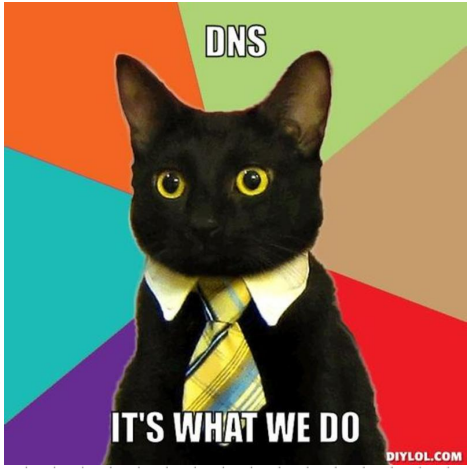

Result

Happy Devving, Happy Opsing

- Using all these techniques to “hand off” unknown to SAAS services we were able to drastically reduce the alerts in our system.
- We no longer have frustration that only 10% of our time can go into development.
- Chaos Monkey is welcome, fully ephemeral.

Questions?

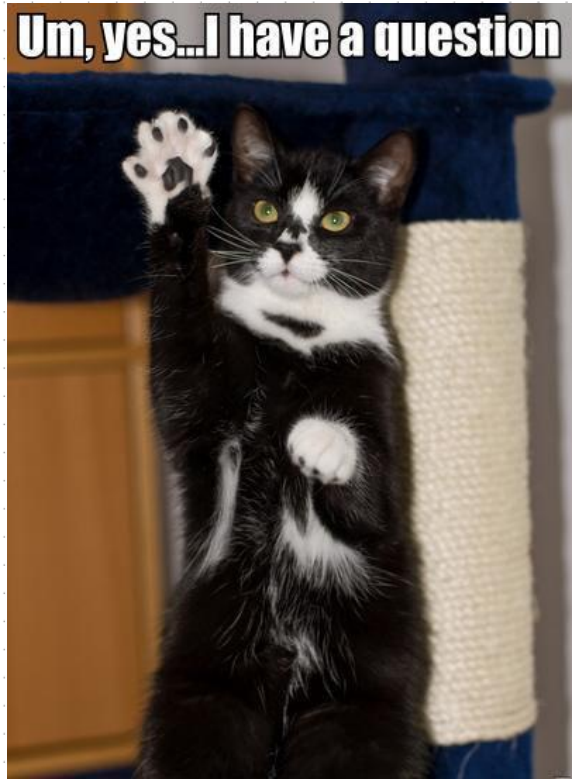

Sprint: Friday

Sprint with the Community on Friday.

We have tasks for every skillset.

Mentors are available for new contributors.

An optional Friday morning workshop for first-time sprinters will help you get set up.

Follow @drupalmentoring.

https://www.flickr.com/photos/amazeelabs/9965814443/in/faves-38914559@N03/


BARCELONA
DRUPALCON 2015
What Did You Think?
Evaluate This Session
barcelona2015.drupal.org/schedule
Thank you!