

# How Netflix Leverages Multiple Regions to Increase Availability

Ruslan Meshenberg

10 März 2014

Please evaluate  
my talk via the  
mobile app!



# Netflix



# Who am I?



- Director, Platform Engineering
- Led Multi-Regional Resiliency
- Leading NetflixOSS Program
- @rusmeshenberg



# Failure

fail·ure

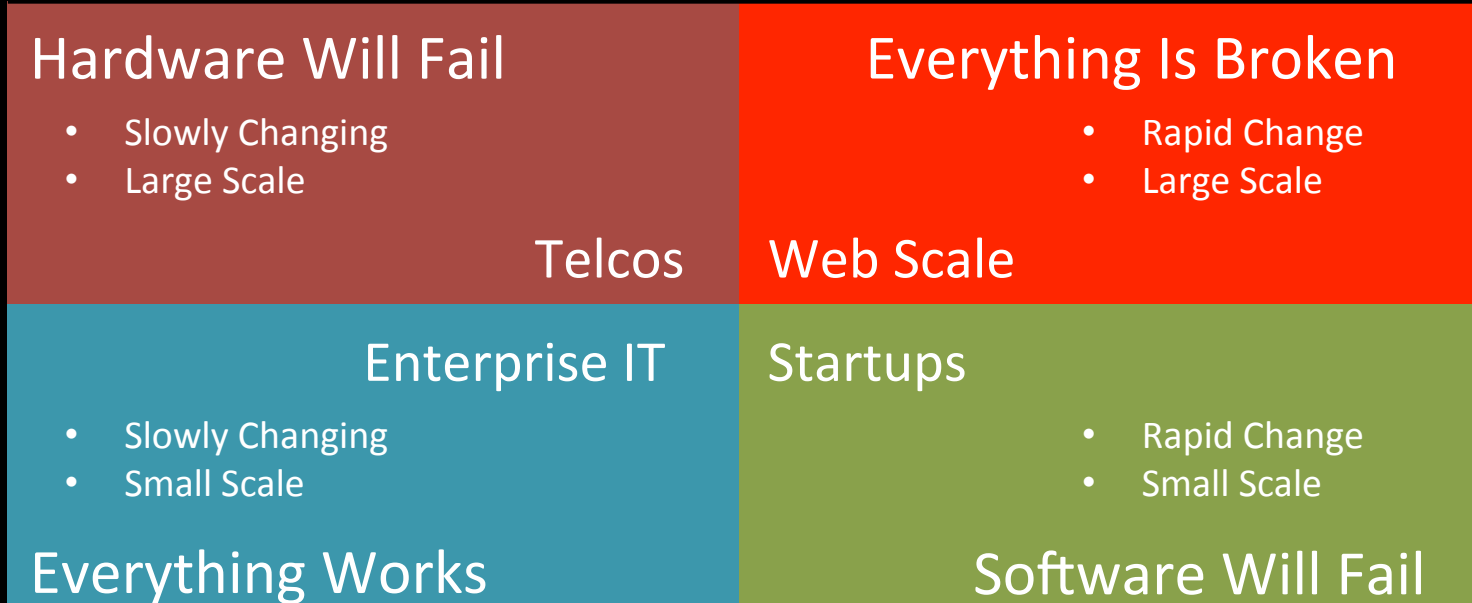
/'fālyər/ 

*noun*

1. lack of success  
"an economic policy that is doomed to failure"  
*synonyms:* lack of success, nonfulfillment, defeat, collapse, foundering
2. the omission of expected or required action.  
"their failure to comply with the basic rules"



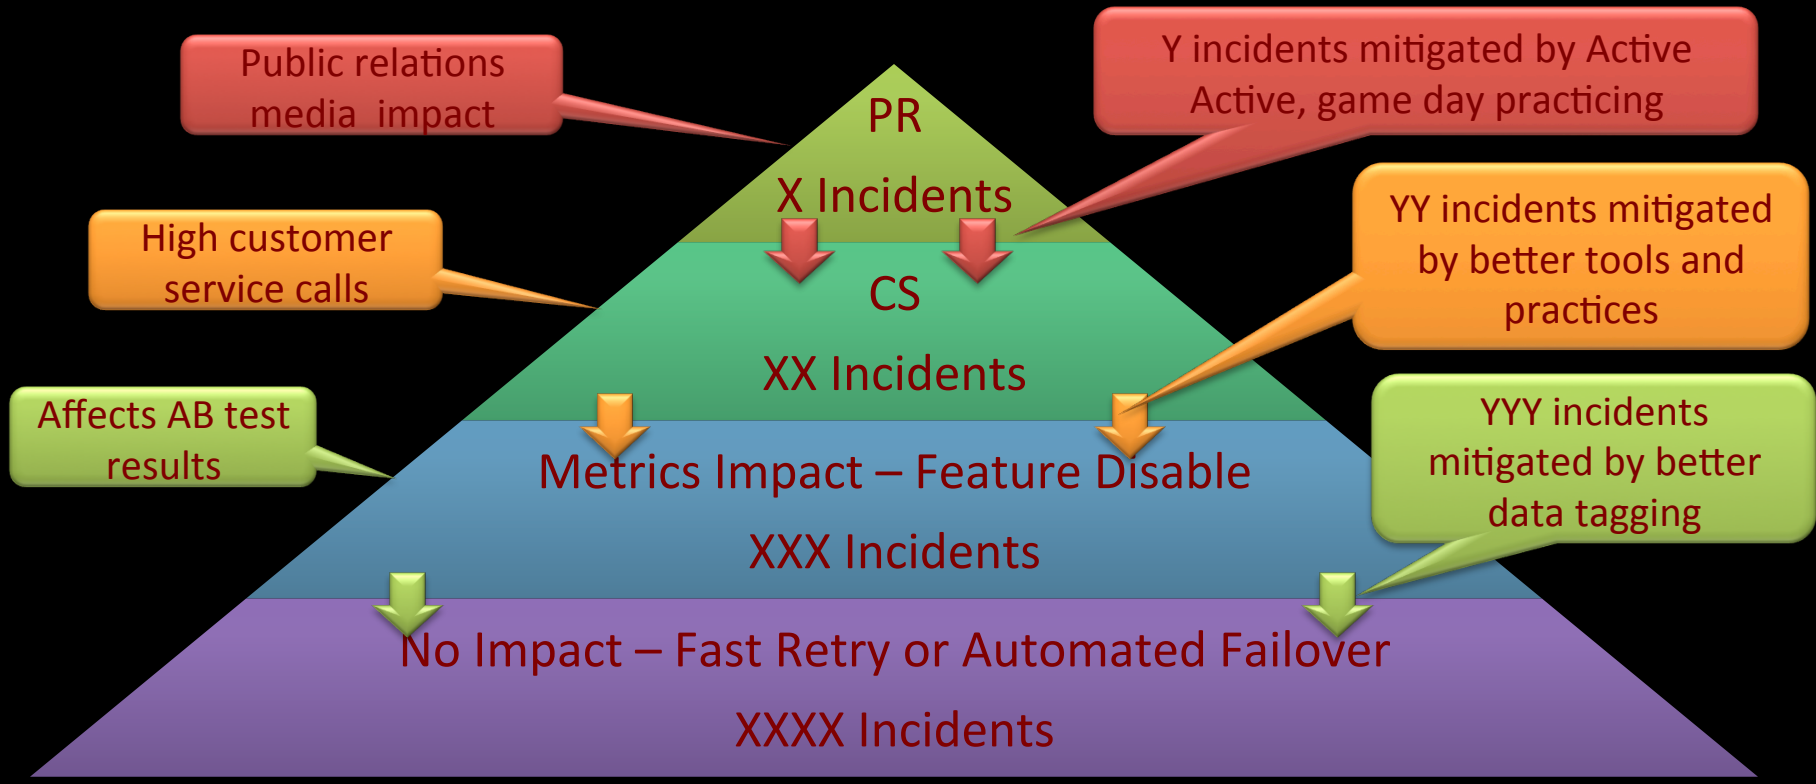
# Assumptions



Scale

Speed

# Incidents – Impact and Mitigation



# Does an Instance Fail?

- It can, plan for it
- Bad code / configuration pushes
- Latent issues
- Hardware failure
- Test with Chaos Monkey



# Does a Zone Fail?

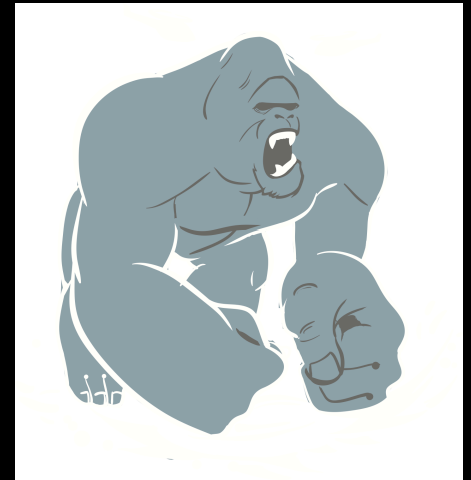
- Rarely, but happened before
- Routing issues
- DC-specific issues
- App-specific issues within a zone
- Test with Chaos Gorilla



# Does a Region Fail?

- Full region – unlikely, very rare
- Individual Services can fail region-wide
- Most likely, a region-wide configuration issue

- Test with Chaos Kong



# Everything Fails... Eventually

- Keep your services running by embracing isolation and redundancy
- Construct a highly agile and highly available service from ephemeral and assumed broken components



# Isolation

- Changes in one region should not affect others
- Regional outage should not affect others
- Network partitioning between regions should not affect functionality / operations

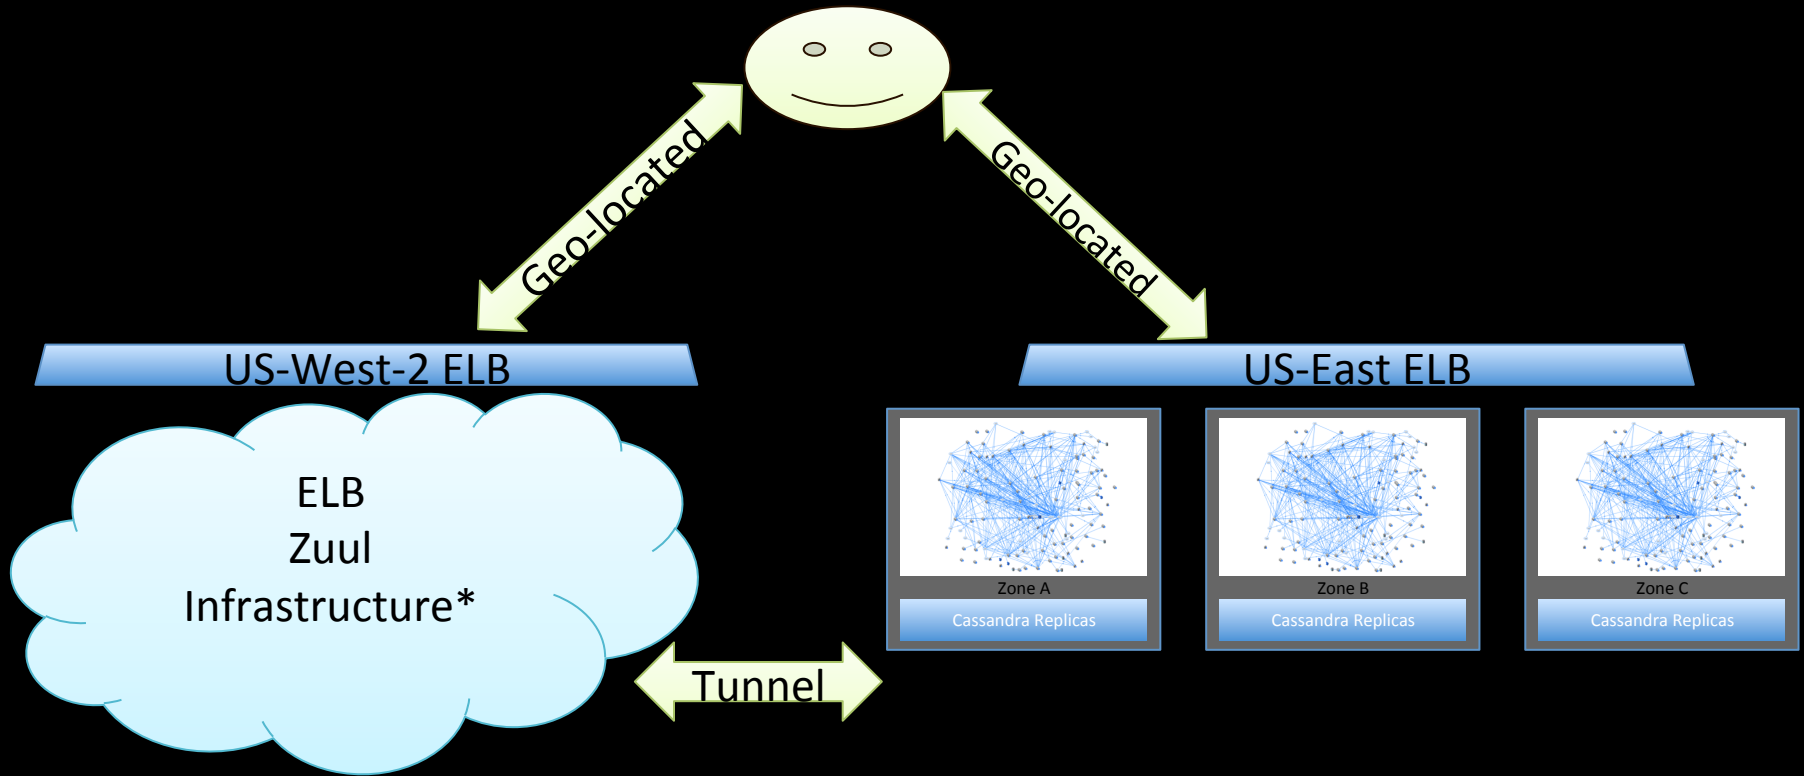
# Redundancy

- Make more than one (of pretty much everything)
- Specifically, distribute services across Availability Zones and regions

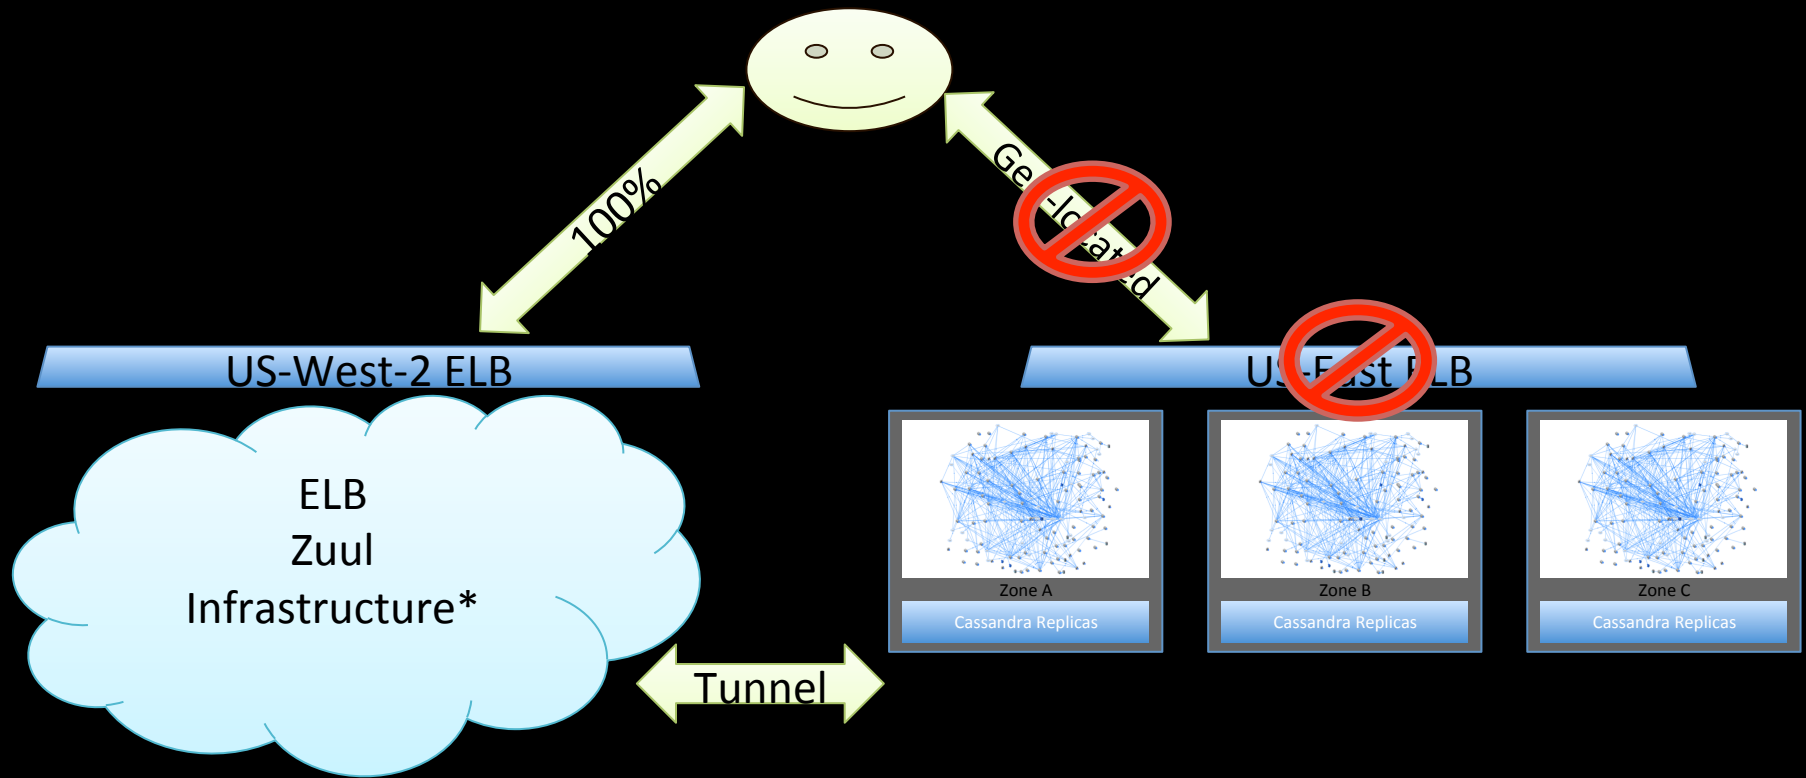
# History: X-mas Eve 2012

- Netflix multi-hour outage
- US-East1 regional Elastic Load Balancing issue
- *“...data was deleted by a maintenance process that was inadvertently run against the production ELB state data”*

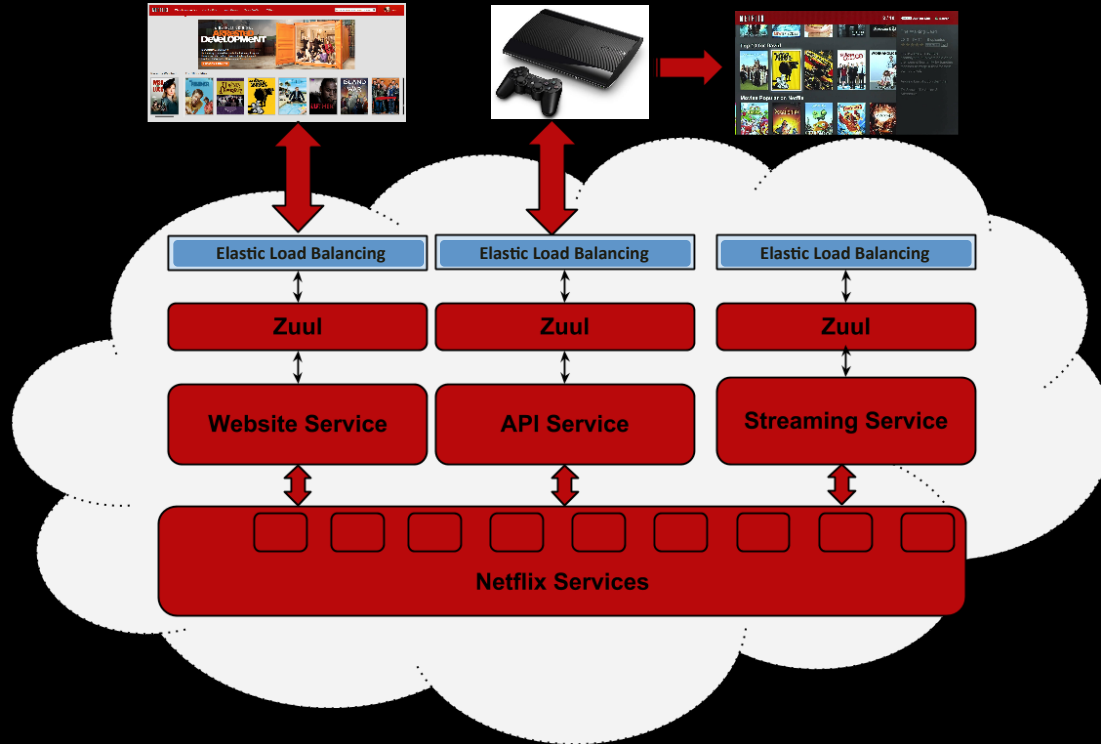
# Isthmus – Normal Operation



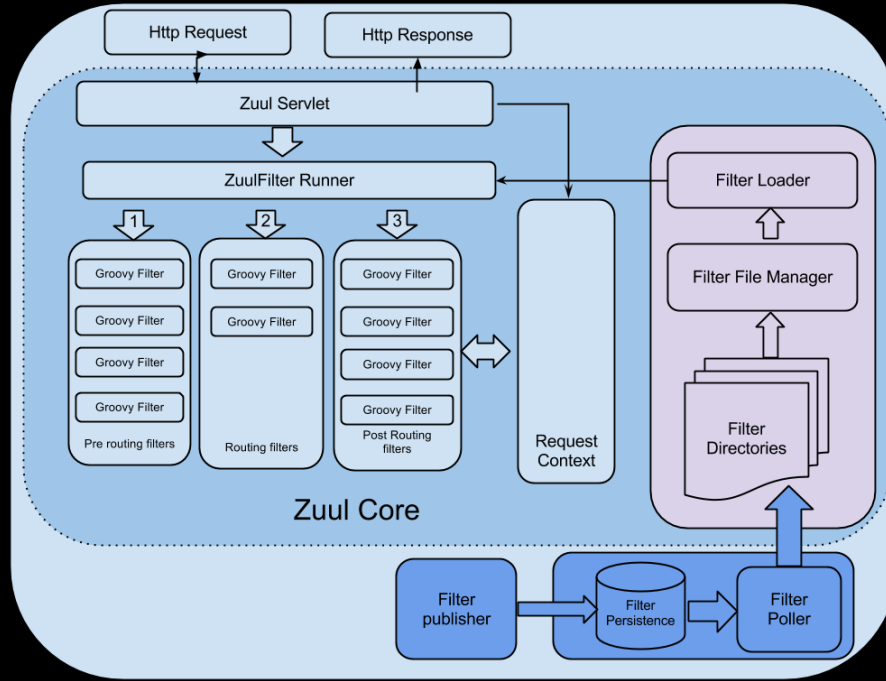
# Isthmus – Failover



# Zuul Overview

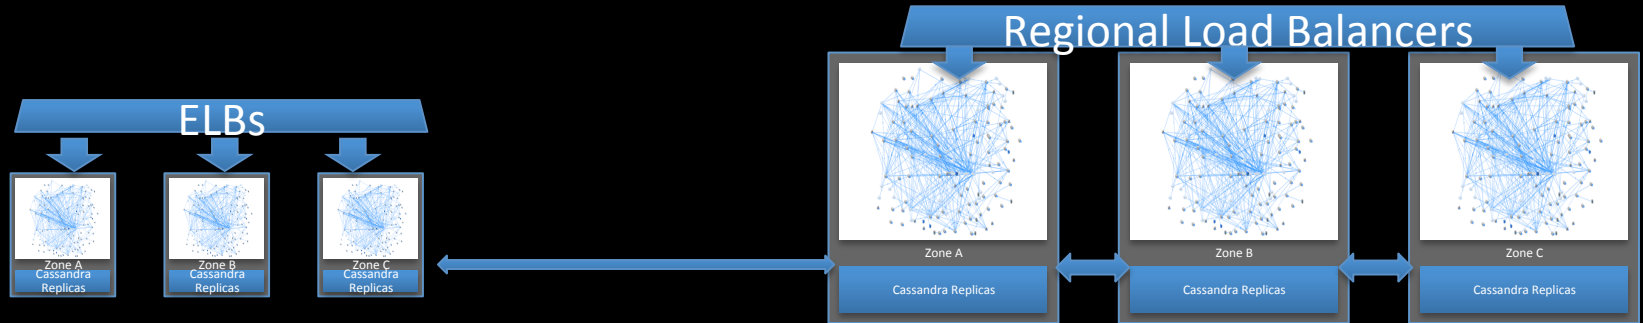


# Zuul





# Denominator – Abstracting the DNS Layer

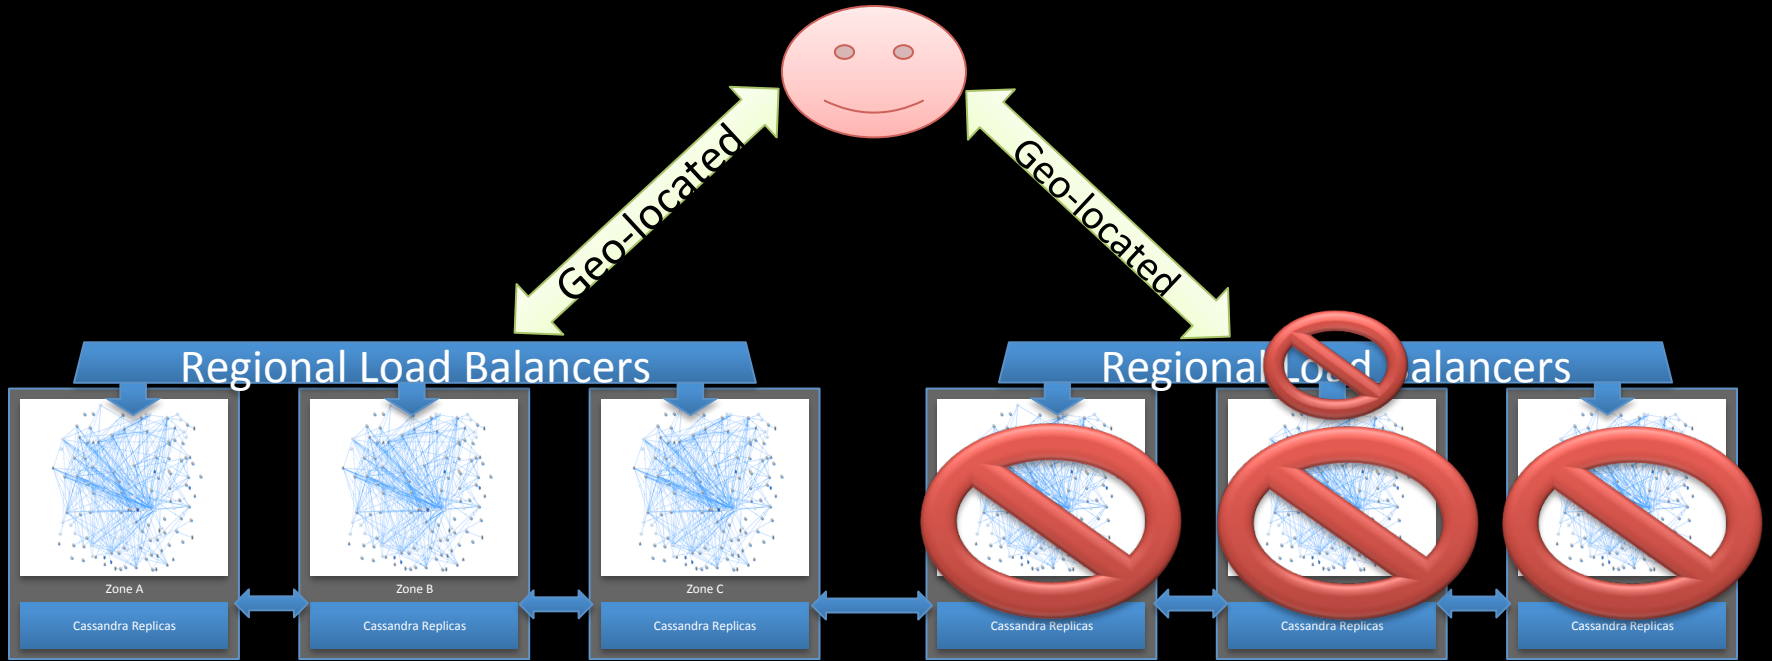


# Isthmus – Only for ELB Failures

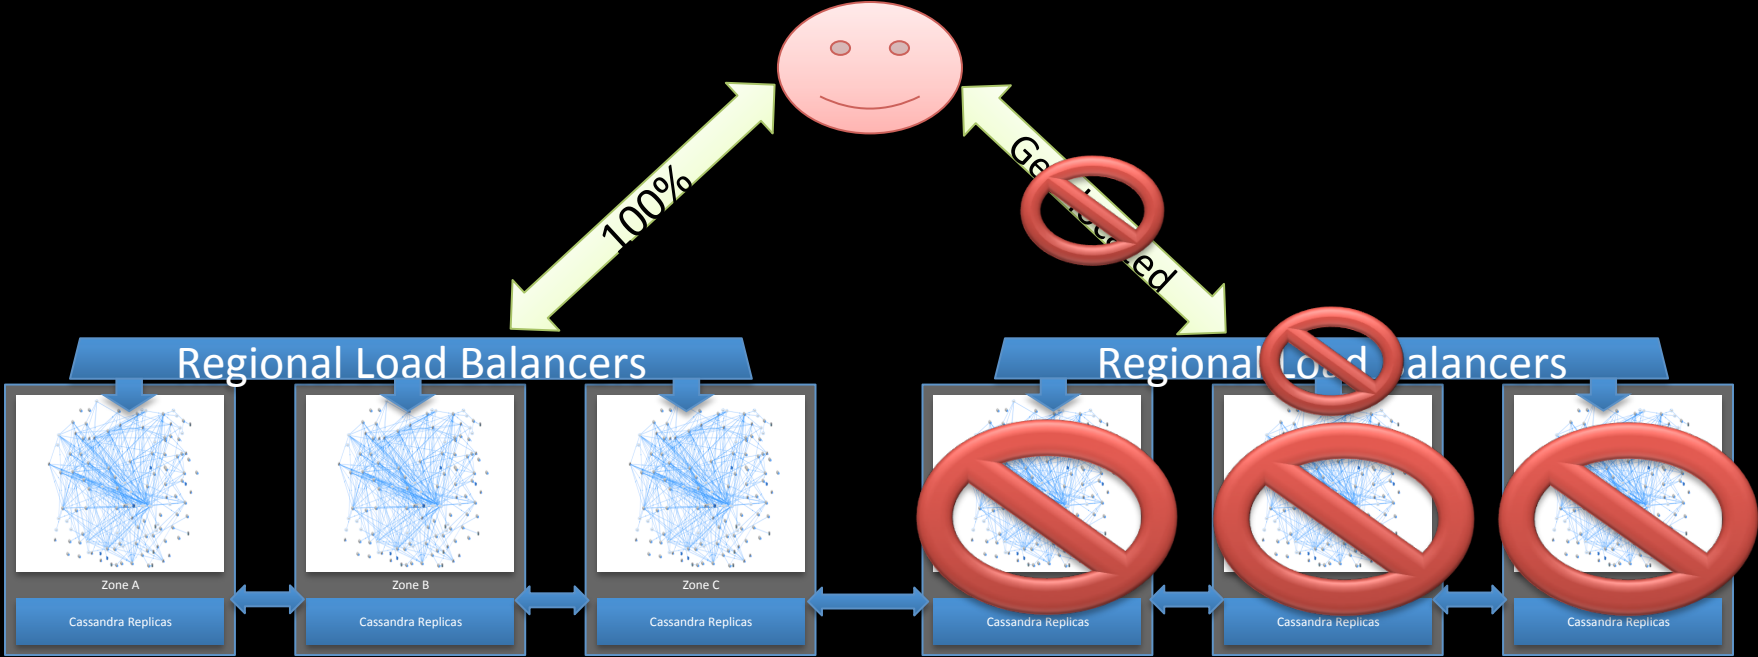
- Other services may fail region-wide
- Not worthwhile to develop one-offs for each one



# Active-Active – Full Regional Resiliency



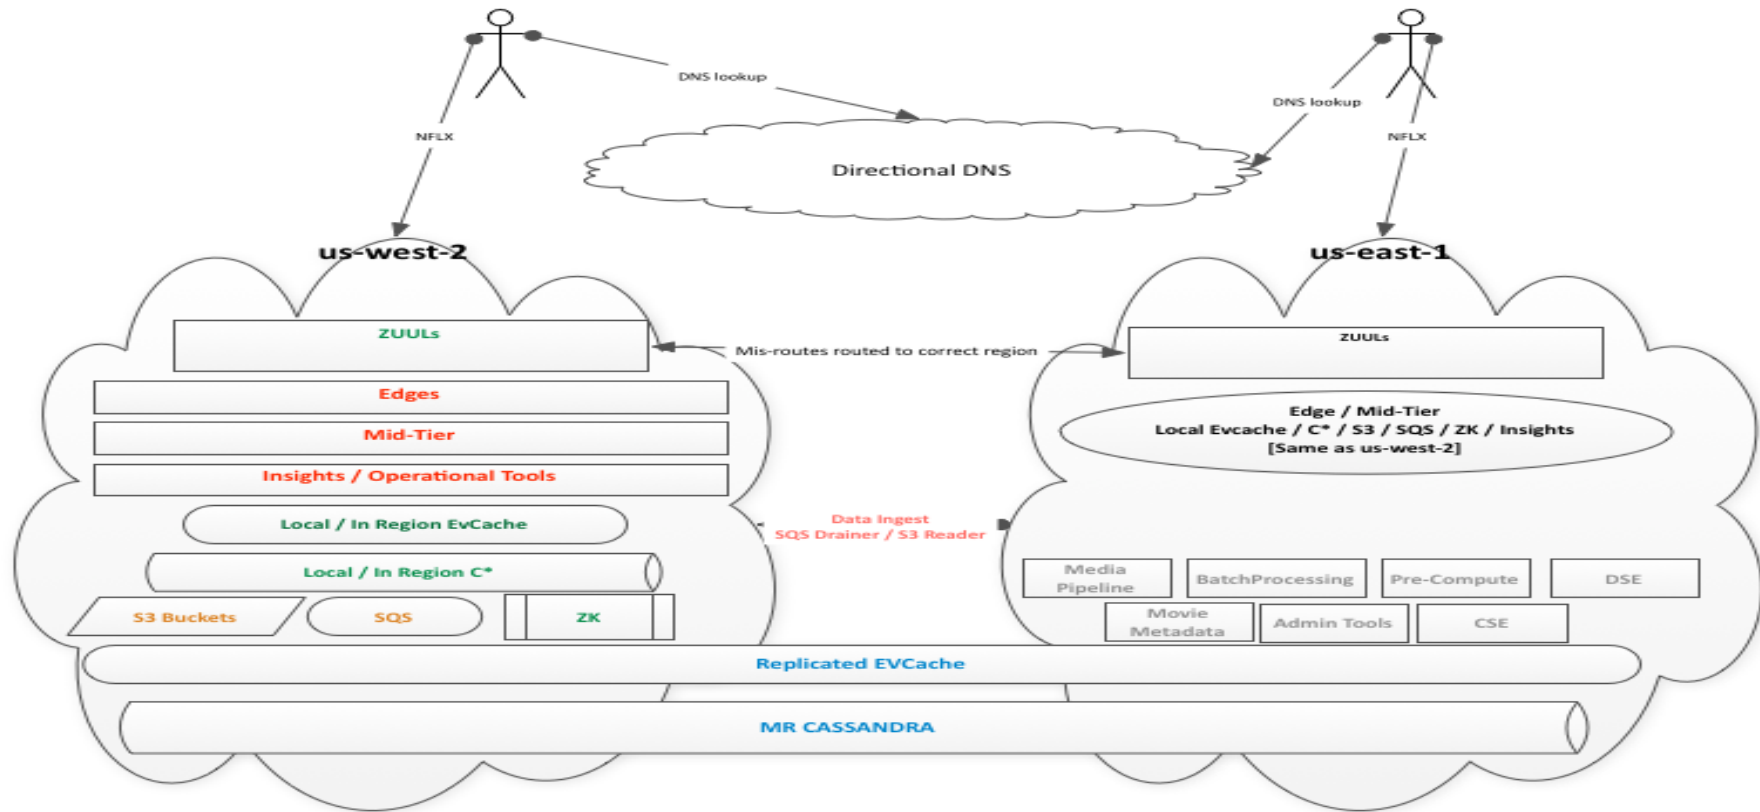
# Active-Active – Failover



Can't we just deploy in 2 regions?



# Active-Active Architecture



# 2 main challenges

- Routing the users to the services
- Replicating the data



# Separating the Data – Eventual Consistency

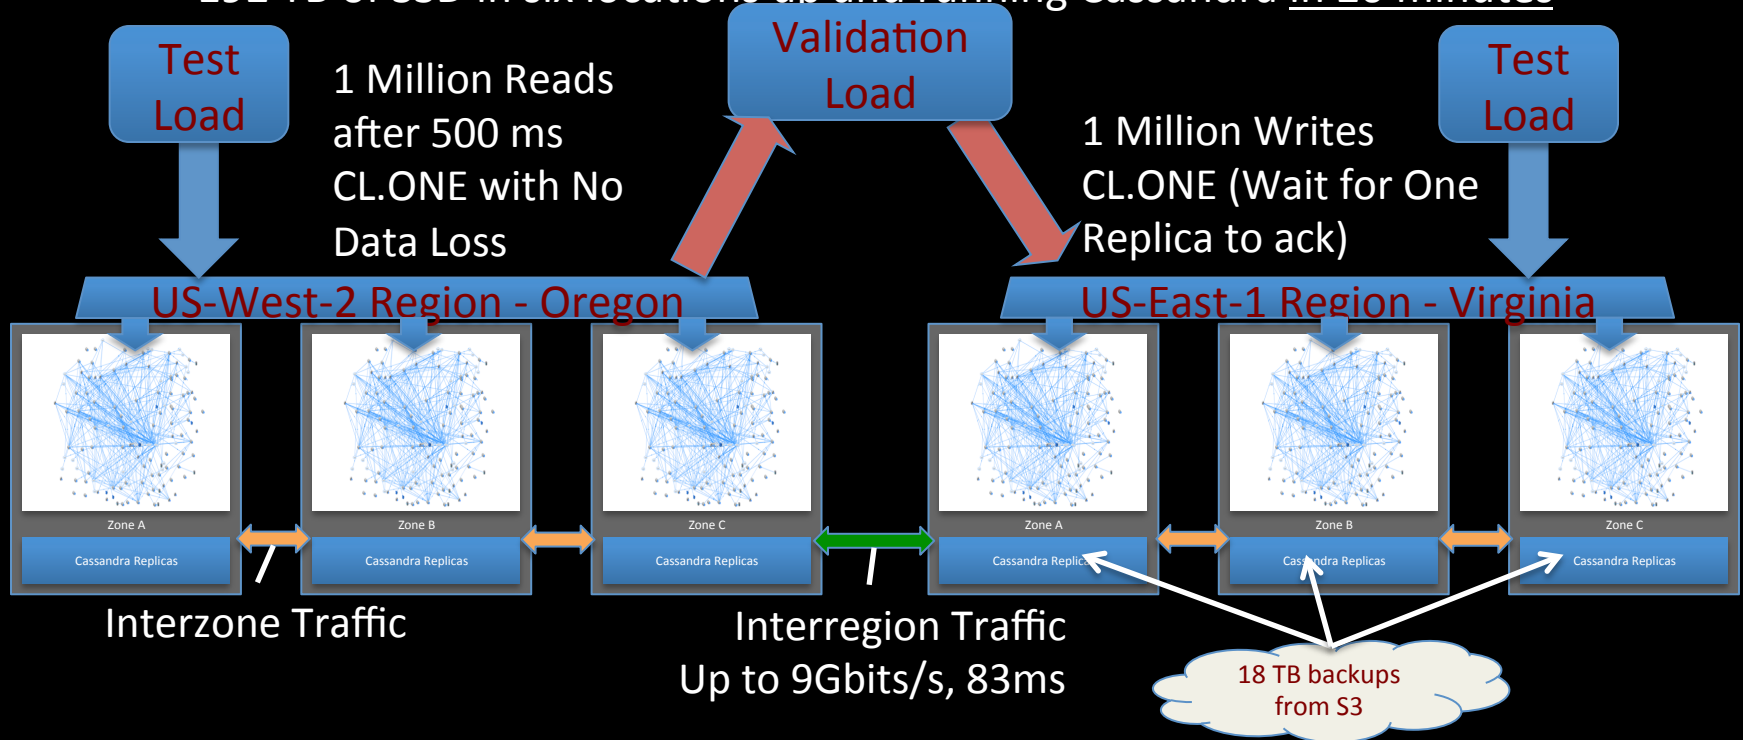
- 2–4 region Cassandra clusters
- Eventual consistency != hopeful consistency

# Benchmarking Global Cassandra

Write intensive test of cross-region replication capacity

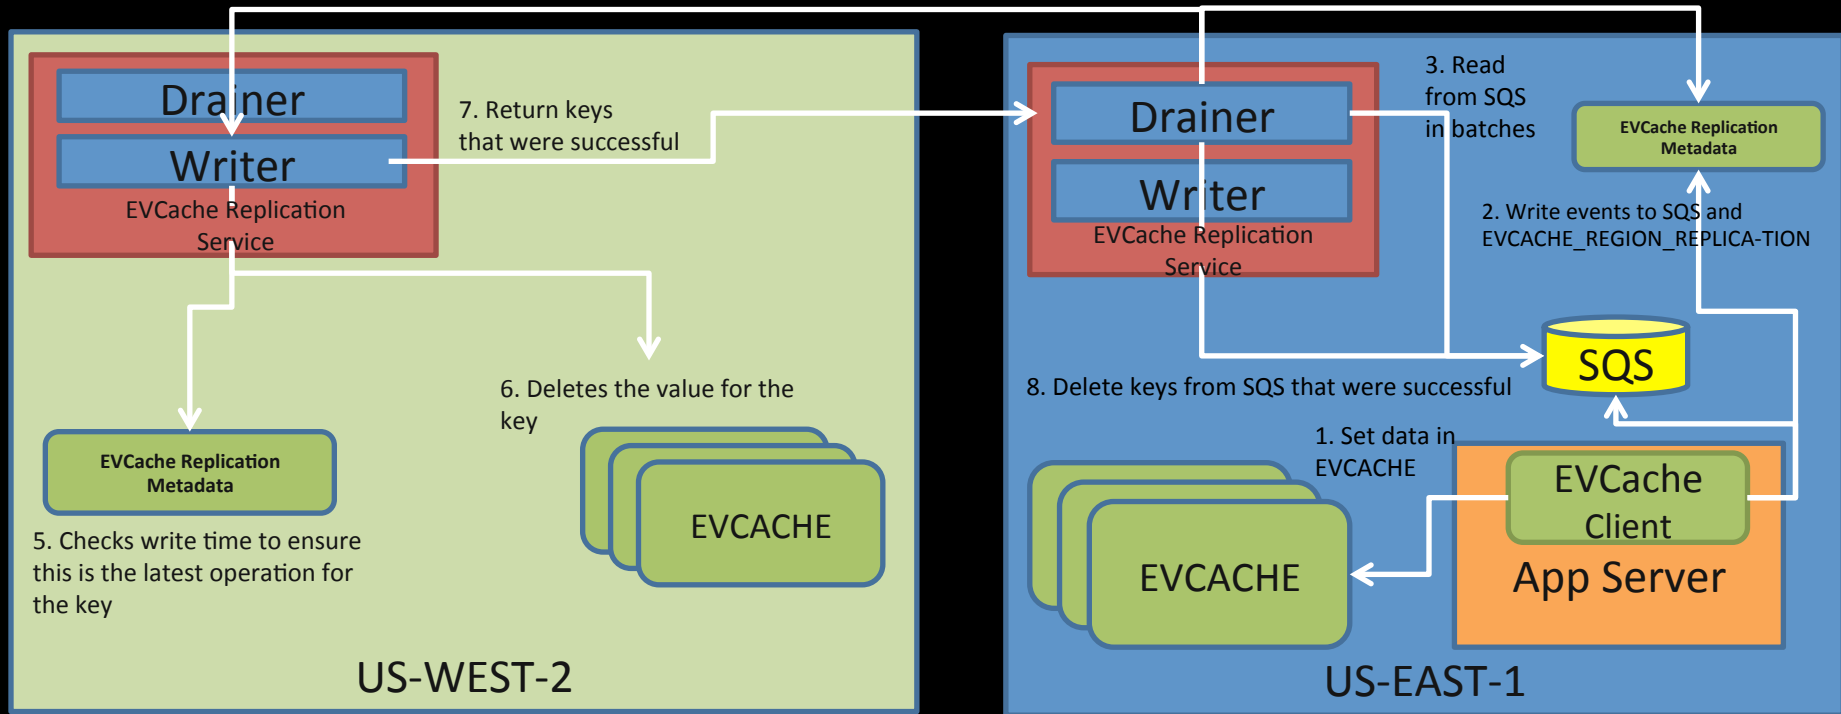
16 x hi1.4xlarge SSD nodes per zone = 96 total

192 TB of SSD in six locations up and running Cassandra in 20 minutes



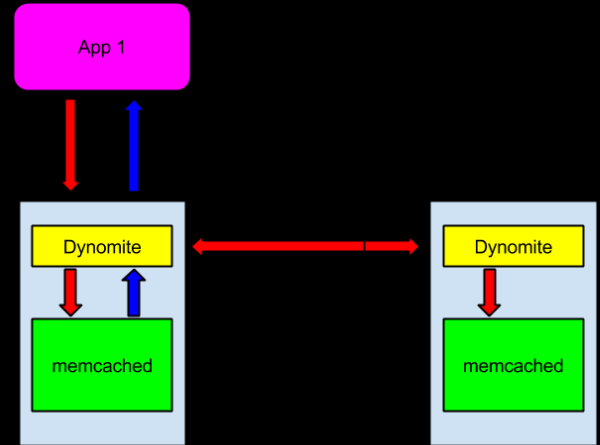
# Propagating EVCache Invalidation

4. Calls Writer with Key, Write Time, TTL & Value after checking if this is the latest event for the key in the current batch.  
Goes cross-region through ELB over HTTPS



# Announcing Dinomyte

- Cross AZ & Region replication to existing Key Value stores
  - Memcached
  - Redis
- Thin Dynamo implementation for replication
- Keep native protocol
  - No code refactoring
- OSS soon!



# Archaius – Region-isolated Configuration

The image shows a web interface for configuring Archaius. On the left, a form titled "Create New Fast Property in te" is partially visible. It includes fields for Name (prop.foo), Value (97), Constraints (int:0-99), Description (property to control %), UpdatedBy (stonse@netflix.com), Notification Email, CMC, and Scope (Not Selected). Overlaid on this is a "Scope Selection" dialog box. The dialog has a title bar with a close button (x). It contains several rows of configuration options, each with a dropdown menu, a count, a filter icon, and a filter input field:

- Region: us-east-1 (Virginia)
- Application: helloworld (1 instances, filter: helloworld)
- Stack: (empty)
- Cluster: helloworld-example (12 instances, filter: filter)
- ASG: helloworld-example-v1 40 (14 instances, filter: filter)

At the bottom of the dialog, there are four buttons: "Select ASG" (blue), "Skip" (orange), "Done" (blue), and "Repeat Selection" (blue with a refresh icon). On the right side of the dialog, there is a "Quick Impact View" section showing the current configuration:

**Quick Impact View**  
Total instances: 2  
Current Scope

region	us-east-1
appld	helloworld
cluster	helloworld-example

# Running Isthmus and Active-Active

# Multiregional Monkeys

- Detect failure to deploy
- Differences in configuration
- Resource differences



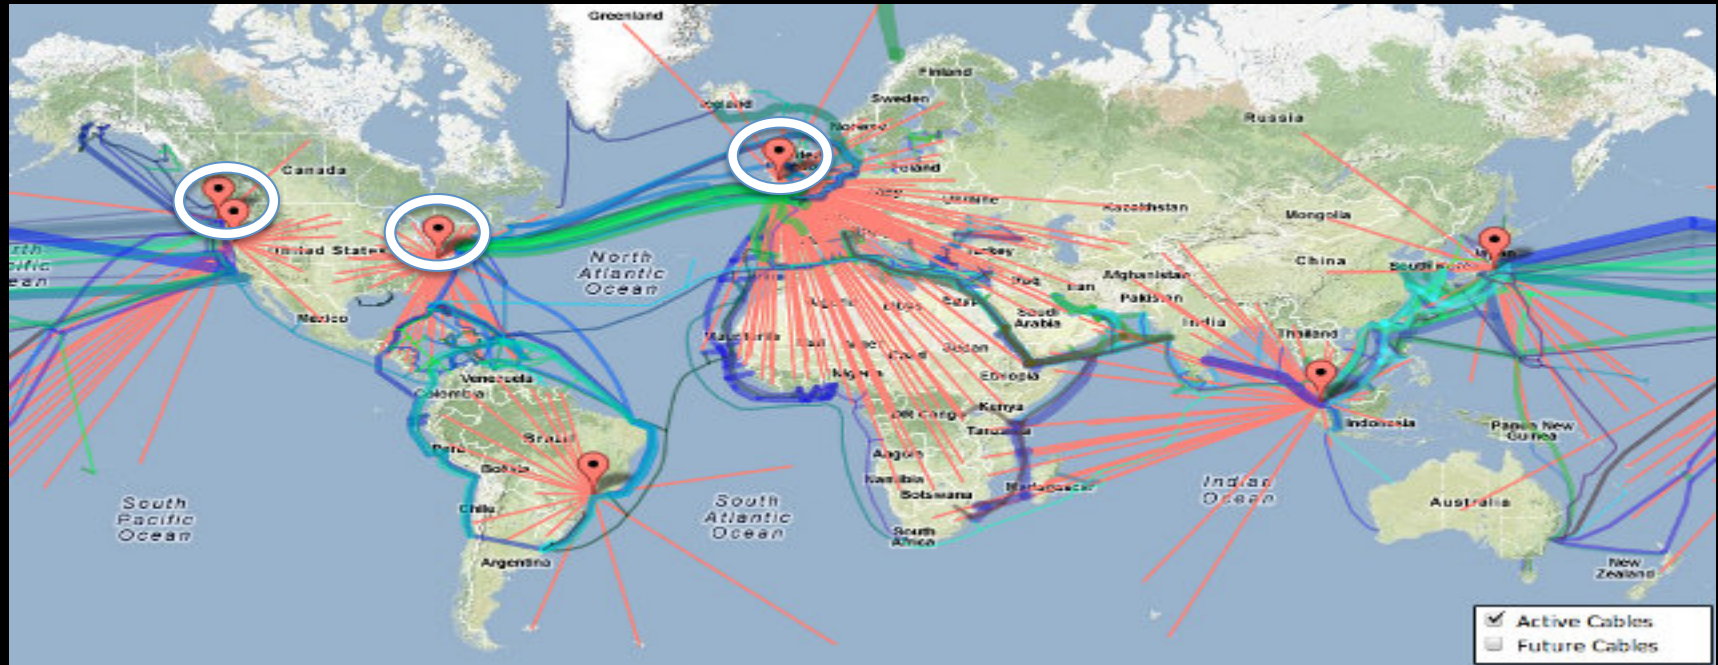


# Operating in N Regions

- Best practices: avoiding peak times for deployment
- Early problem detection / rollbacks
- Automated canaries / continuous delivery



# 6 deployments



At least 36 steps

# Automating deployment

The screenshot displays the Mimir console interface for a service named 'platformservice'. The navigation bar includes 'Mimir', 'Builds', 'Deployments', 'AMs', 'Clusters', 'ASGs', and 'Workflows'. The main content area is divided into sections for 'LATEST BUILD' and 'TEST'.

**LATEST BUILD**

Buttons: [Promote to Test](#)

Jenkins Job: WE-WAPP-platformservice  
Status: Pass  
Last Built: 659  
Build Time: 09/27/2013

**TEST**

Buttons: [Promote to Production](#)

**TEST | US-WEST-2 | PLATFORMSERVICE-INT**

ASG Name	Build	Status	Instances	Created
platformservice-int-v009	658	<span style="color: red;">Down</span>	2	09/20/2013
platformservice-int-v010	659	<span style="color: green;">Up</span>	2	09/27/2013

**TEST | US-EAST-1 | PLATFORMSERVICE**

ASG Name	Build	Status	Instances	Created
platformservice-v105	659	<span style="color: green;">Up</span>	6	09/27/2013

**TEST | US-WEST-2 | PLATFORMSERVICE**

ASG Name	Build	Status	Instances	Created
platformservice-v047	658	<span style="color: red;">Down</span>	2	09/20/2013
platformservice-v048	659	<span style="color: green;">Up</span>	2	09/27/2013

**TEST | EU-WEST-1 | PLATFORMSERVICE-INT**

ASG Name	Build	Status	Instances	Created
platformservice-int-v011	658	<span style="color: red;">Down</span>	2	09/20/2013
platformservice-int-v012	659	<span style="color: green;">Up</span>	2	09/27/2013

**TEST | US-EAST-1 | PLATFORMSERVICE-INT**

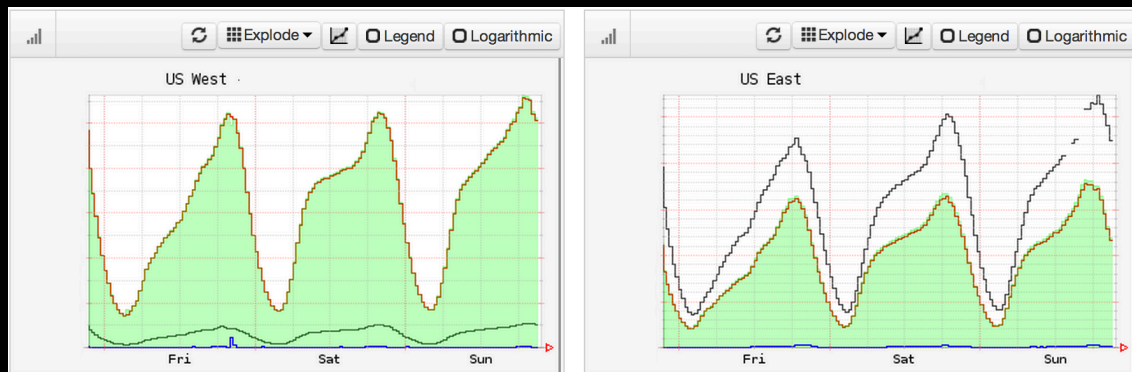
ASG Name	Build	Status	Instances	Created
platformservice-int-v016	658	<span style="color: red;">Down</span>	2	09/20/2013
platformservice-int-v017	659	<span style="color: green;">Up</span>	2	09/27/2013

**TEST | US-WEST-1 | PLATFORMSERVICE**

ASG Name	Build	Status	Instances	Created
----------	-------	--------	-----------	---------

# Monitoring and Alerting

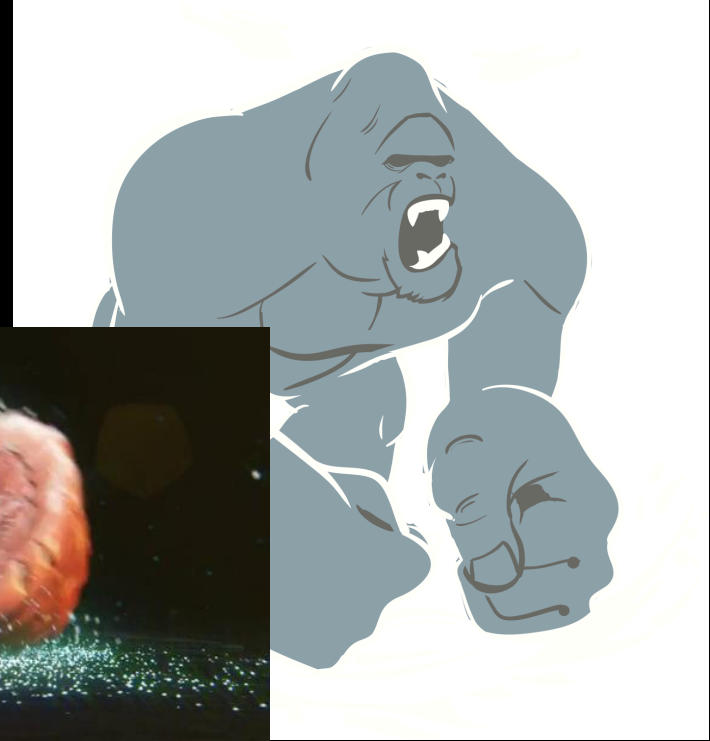
- Per region metrics
- Global aggregation
- Anomaly detection



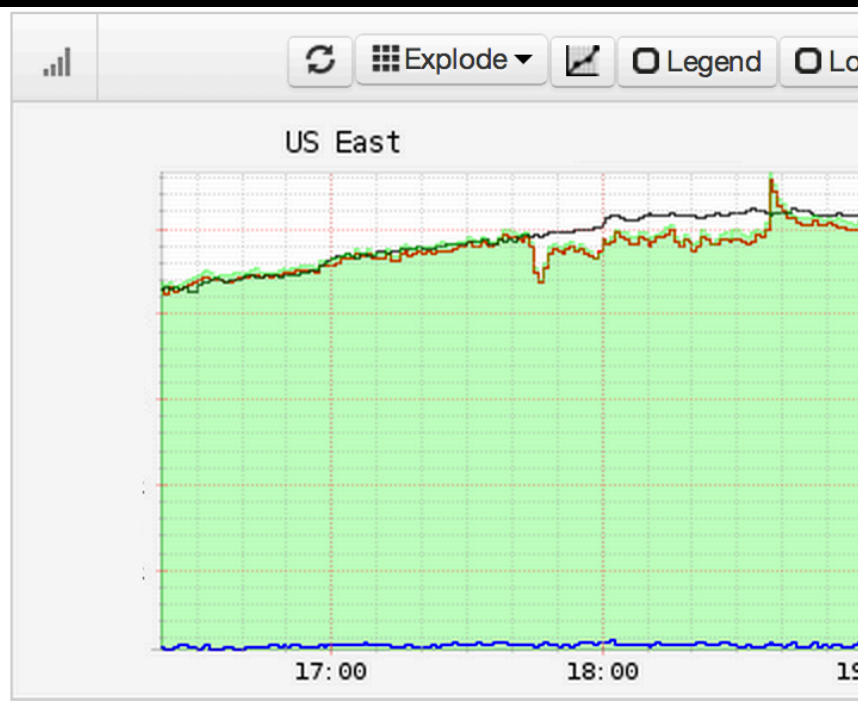
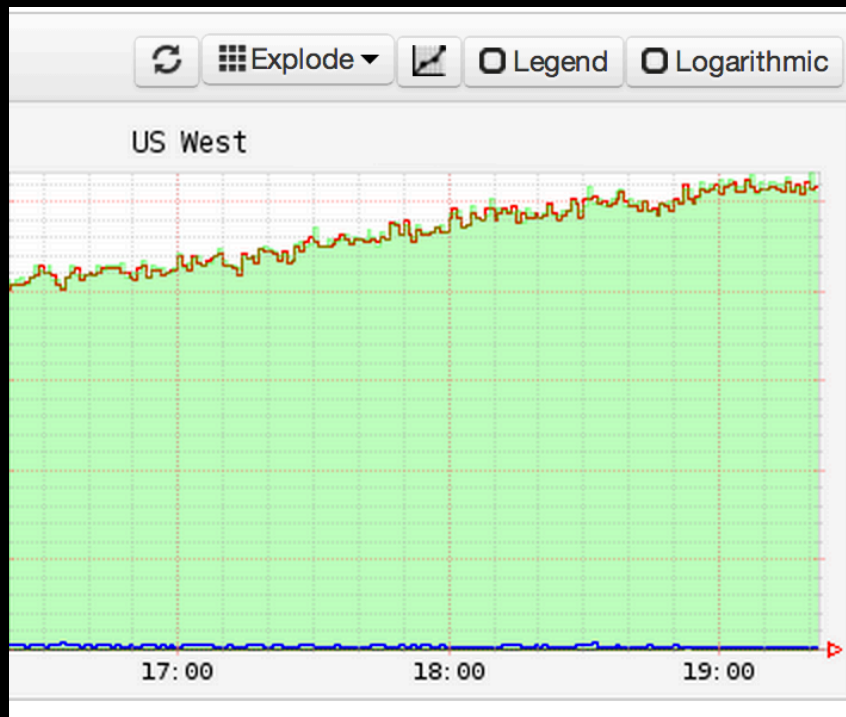
# Failover and Fallback

- Directional DNS changes (Or Route53)
- For fallback, ensure data consistency
- Some challenges
  - Cold cache
  - Autoscaling
- (Iterate, Automate)<sup>+</sup>

# Validating the Whole Thing Works

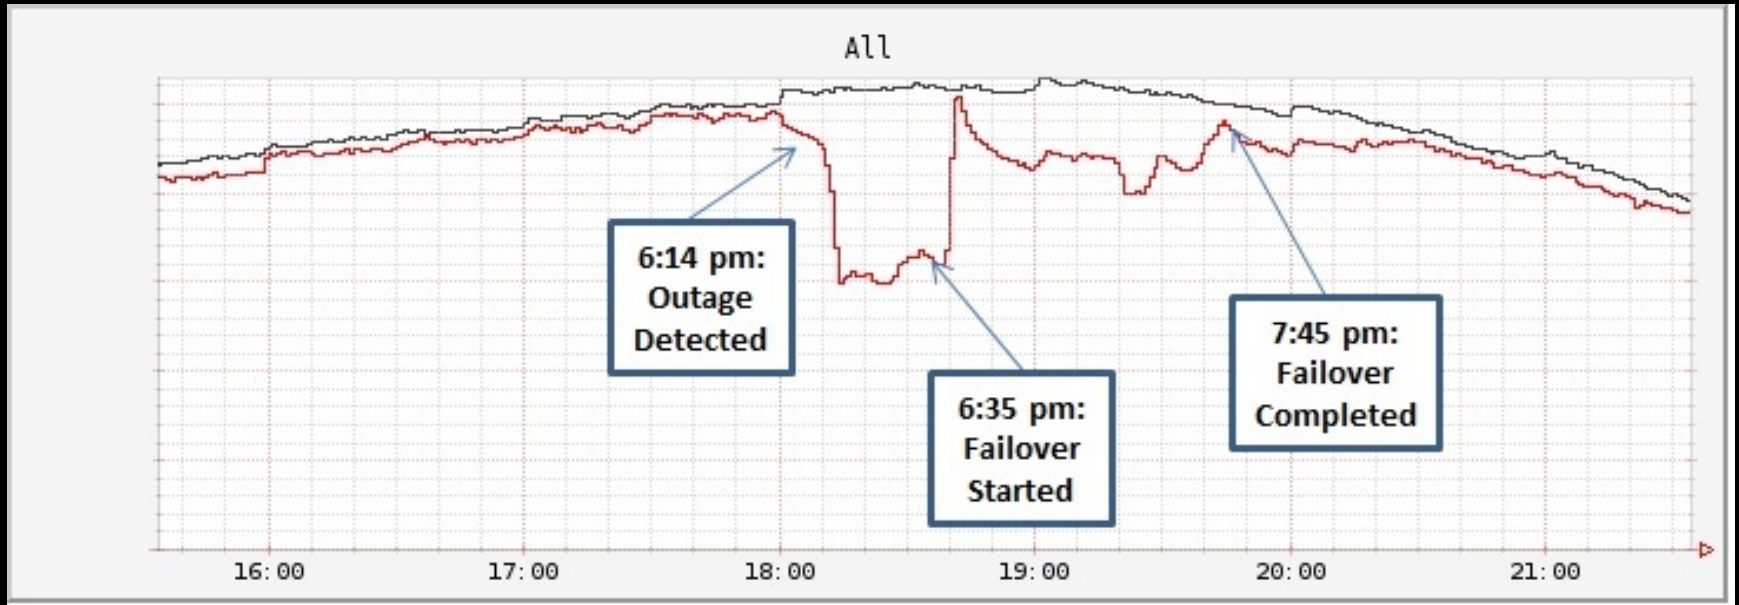


# Does Isolation Work?

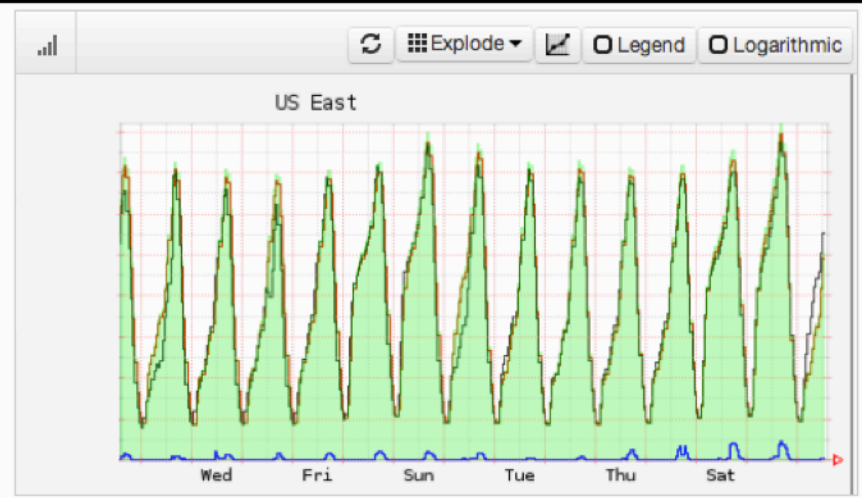
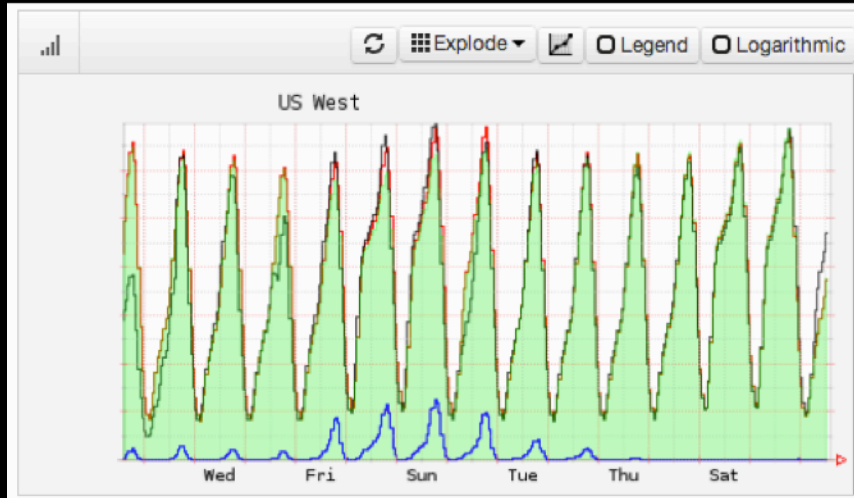




# Does Failover work?



# Steady State



# Wii.FM?

The screenshot displays the Netflix OSS Open Source Software Center interface. At the top, it features the Netflix logo, the OSS logo, and the text "Netflix Open Source Software Center". Below this, there are navigation tabs for "Repositories", "Current Timeline", and "Powered By Netflix OSS".

The main content area is organized into several sections:

- Repositories:** This section contains three icons: "MYSTEX" (a stylized eye), "SIRIANNONY" (a man's face), and "TURJINI" (a person's face).
- Cloud Management Applications:** This section contains two icons: "ICE" (a cartoon character) and "ASGARD" (a hammer).
- Perisatence Systems:** This section contains a row of seven icons: "ASTYMAX" (a dark scene), "CASLMETER" (a person's face), "CURATOR" (a gear), "EVAOAE" (a red 'E'), "EXHIBITOR" (a person's face), "PREAM" (a person's face), and "STAIGH" (a person's face).
- Platform Libraries:** This section contains a row of ten icons: "ANCHOUS" (a landscape), "GENOMINATOR" (a cartoon character), "FESON" (a person's face), "KARYON" (a dog), "REBORN" (a red bow), "SERVO" (a green circle), "BLITZAI" (a person's face), and "GOVERNATOR" (a person's face).
- Infrastructure Services:** This section contains a row of four icons: "EDA" (a woman's face), "SUPO" (a person's face), "EUREKA" (a person's face), and "ZIAL" (a dog's face).

**NETFLIX**

**OSS**



We're here to help you get to global scale...  
Apache Licensed Cloud Native OSS Platform

<http://netflix.github.com>









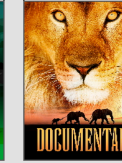






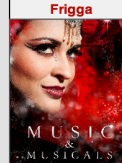

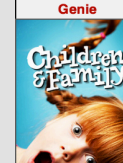

















# Technical Indigestion – what do all these do?

NETFLIX **OSS** Netflix Open Source Center

Repositories Commit Timeline Mailing Lists Community Powered By NetflixOSS Around the Web AMLs

Our Repositories  

35 public repos  
27 members

<b>Aminator</b> 	<b>Archaius</b> 	<b>Asgard</b> 	<b>Astyanax</b> 	<b>Blitz4j</b> 	<b>Brutal</b> 	<b>CassJMeter</b> 	<b>Cloud-Prize</b> 	<b>Curator</b> 	<b>Denominator</b> 	<b>Edda</b> 	<b>Eureka</b> 
<b>EVCache</b> 	<b>Exhibitor</b> 	<b>Feign</b> 	<b>Frigga</b> 	<b>Gcviz</b> 	<b>Genie</b> 	<b>Glisten</b> 	<b>Governator</b> 	<b>Hystrix</b> 	<b>Ice</b> 	<b>Karyon</b> 	<b>Lipstick</b> 
<b>Netflix-Graph</b> 	<b>NIWebCrypto</b> 	<b>Priam</b> 	<b>Pytheas</b> 	<b>Recipes-Rss</b> 	<b>Ribbon</b> 	<b>RxJava</b> 	<b>Servo</b> 	<b>SimianArmy</b> 	<b>Turbine</b> 	<b>Zuul</b> 	

# Asgard - Manage Red/Black Deployments

The screenshot shows the Asgard web interface for managing a cluster of Sequential Auto Scaling Groups. The URL in the browser is `asgardprod/us-east-1/cluster/show/obiwan`. The page title is "Manage Cluster of Sequential Auto Scaling Groups".

Two Auto Scaling Groups are visible:

- obiwan-v063**: 9 instances grouped by state. A table shows 9 instances in the "InService" state, with a "Build" of 580 and "ELB" status of "OUT\_OF\_SERVICE".
- obiwan-v064**: 9 instances grouped by state. A table shows 9 instances in the "InService" state, with a "Build" of 583 and "ELB" status of "UP".

Annotations in green text indicate traffic distribution:

- "No traffic on old version" points to the "OUT\_OF\_SERVICE" status of the v063 instances.
- "Live traffic on new version" points to the "UP" status of the v064 instances.
- "This cluster contains two ASGs" points to the cluster header.

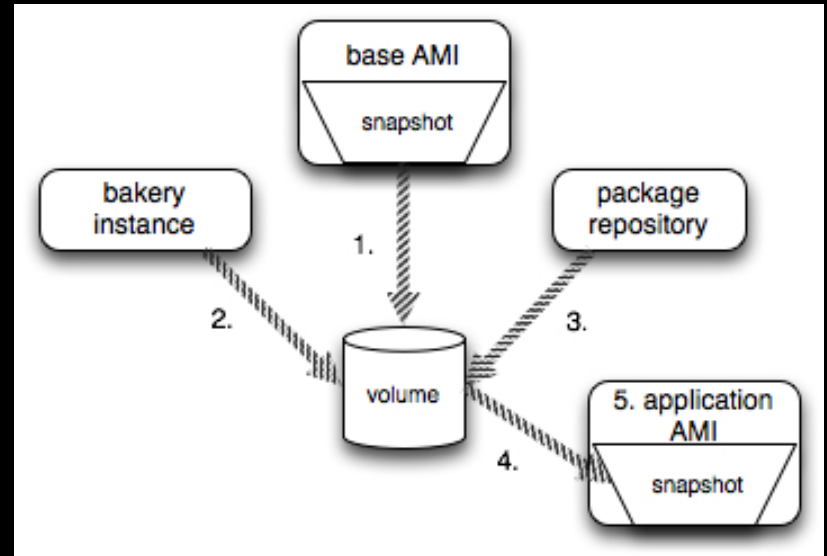
On the right, the "Create Next Group" section shows configuration for "obiwan-v065":

- AMI Image ID: 179123456789/obiwan-41.2-1417305
- Instance Type: m1.large
- Instance Counts: Min: 9, Desired: 9, Max: 12
- After launch:  Wait for Discovery health check pass
- Button: Create Next Group obiwan-v065



# Automatically Baking AMIs with Aminator

- AutoScaleGroup instances should be identical
- Base plus code/config
- Immutable instances
- Works for 1 or 1000...
- Aminator Launch
  - Use Asgard to start AMI or
  - CloudFormation Recipe



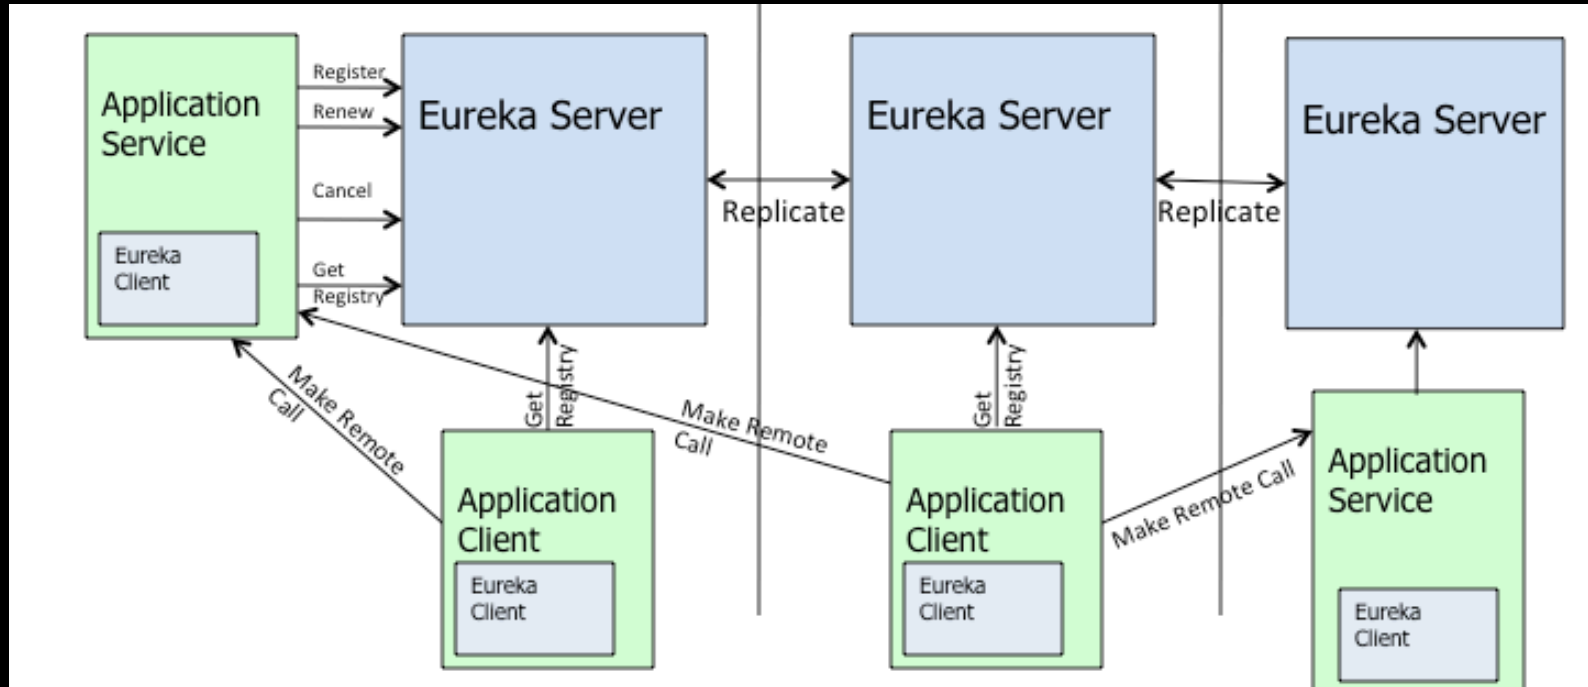


# Discovering your Services - Eureka

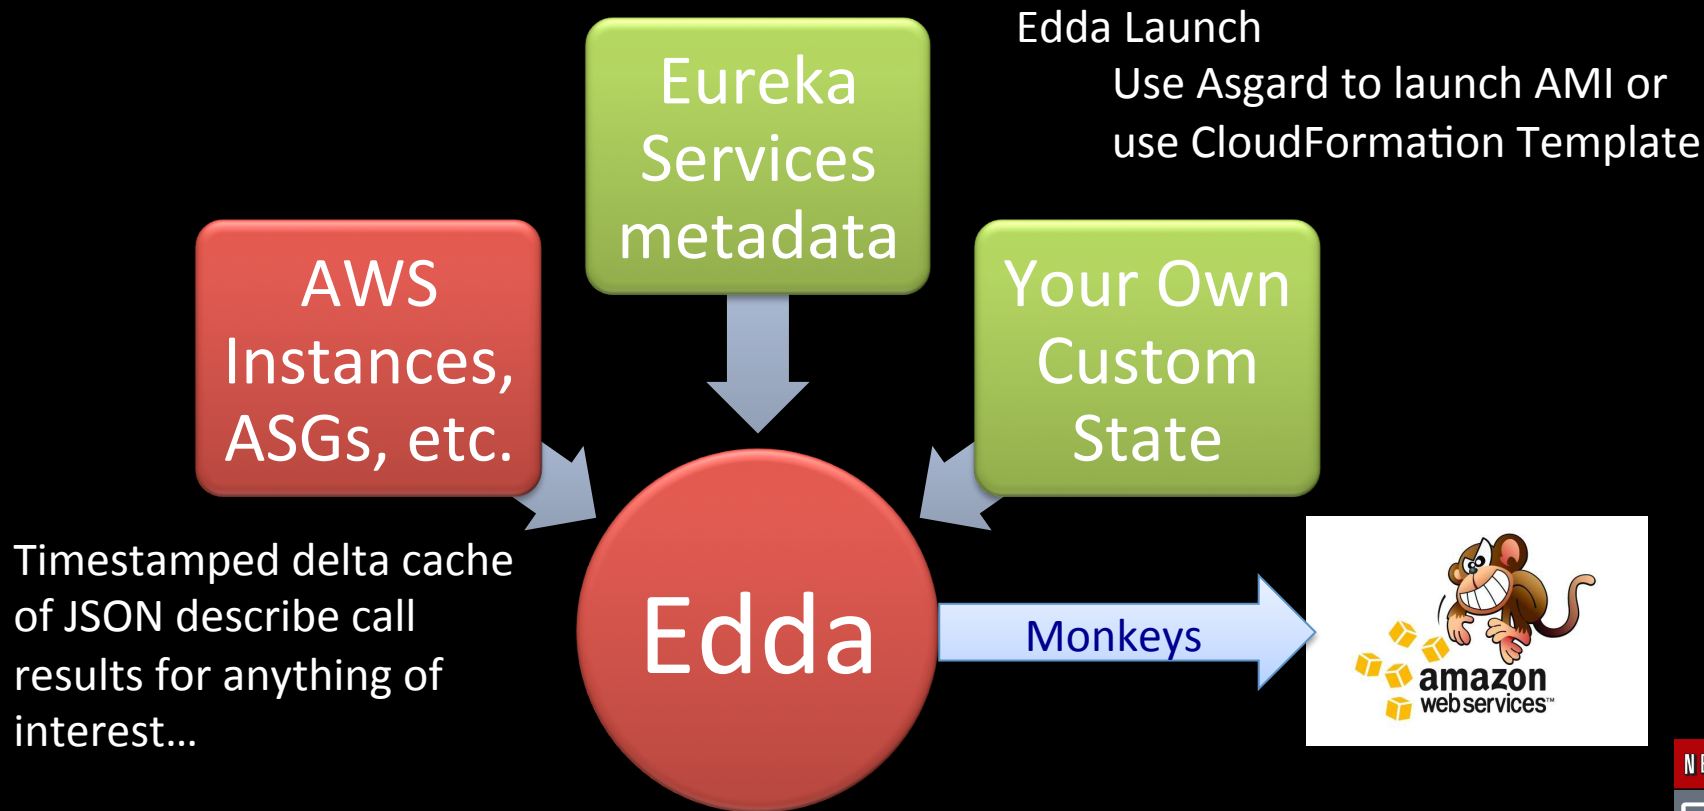
EVCACHE	ami-e794f78e (3), ami-c579ecac (1),	us-east-1c (3), us-east-1d (1),	UP (4) - <a href="#">i-9fd</a> , <a href="#">i-08</a> , <a href="#">i-78e</a> , <a href="#">i-45f</a>
---------	-------------------------------------	---------------------------------	---

- Map applications by name to
  - AMI, instances, Zones
  - IP addresses, URLs, ports
  - Keep track of healthy, unhealthy and initializing instances
- Eureka Launch
  - Use Asgard to launch AMI or use CloudFormation Template

# Deploying Eureka Service – 1 per Zone



# Searchable state history for a Region / Account



# Edda Query Examples

Find any instances that have ever had a specific public IP address

```
$ curl "http://edda/api/v2/view/instances;publicIpAddress=1.2.3.4;_since=0"  
["i-0123456789", "i-012345678a", "i-012345678b"]
```

Show the most recent change to a security group

```
$ curl "http://edda/api/v2/aws/securityGroups/sg-0123456789;_diff;_all;_limit=2"  
--- /api/v2/aws.securityGroups/sg-0123456789;_pp;_at=1351040779810  
+++ /api/v2/aws.securityGroups/sg-0123456789;_pp;_at=1351044093504  
@@ -1,33 +1,33 @@  
{  
...  
  "ipRanges" : [  
    "10.10.1.1/32",  
    "10.10.1.2/32",  
+   "10.10.1.3/32",  
-   "10.10.1.4/32"  
...  
}
```

# Archaius – Property Console

Persisted Properties Console : TEST us-east-1

Properties | Diff | Journal Logs | Documentation | Reports

Home / FastProperty Explorer /

### Create New Fast Property in test

Name:

Value:

Constraints:

Description:

UpdatedBy:

Notification Email:

CMC:

Scope:

#### Scope Selection

Region:

Application:  1

Stack:

Cluster:  12

ASG:  14

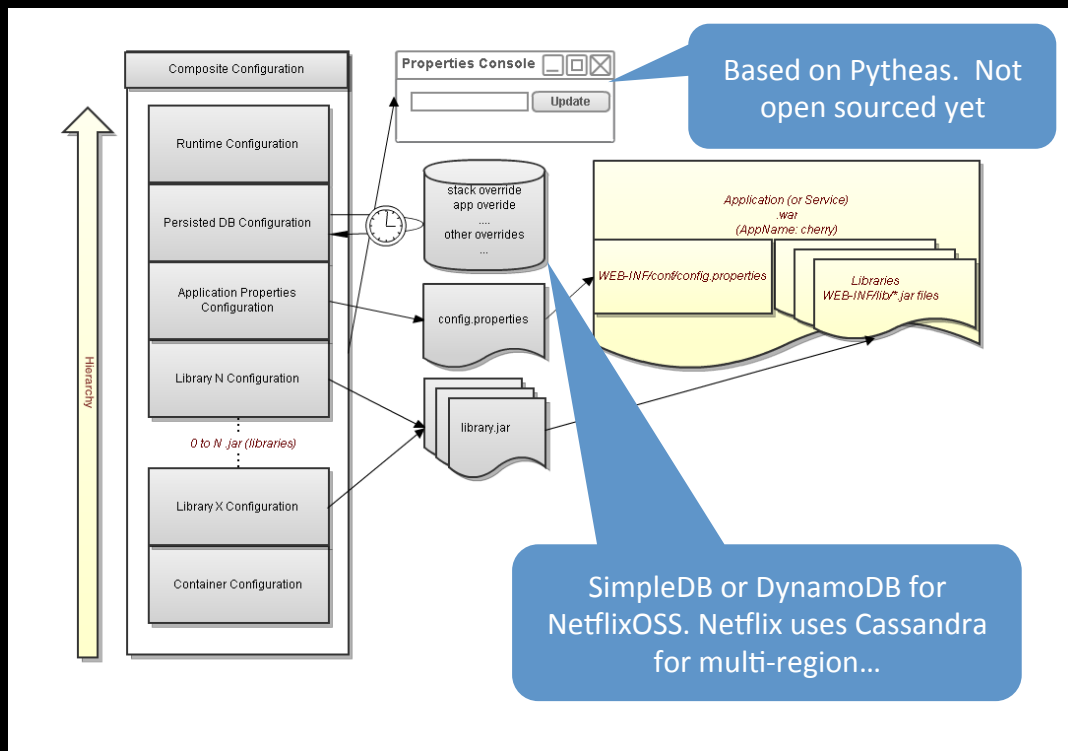
#### Quick Impact View

Total instances: 2

Current Scope

region	us-east-1
appld	helloworld
cluster	helloworld-example

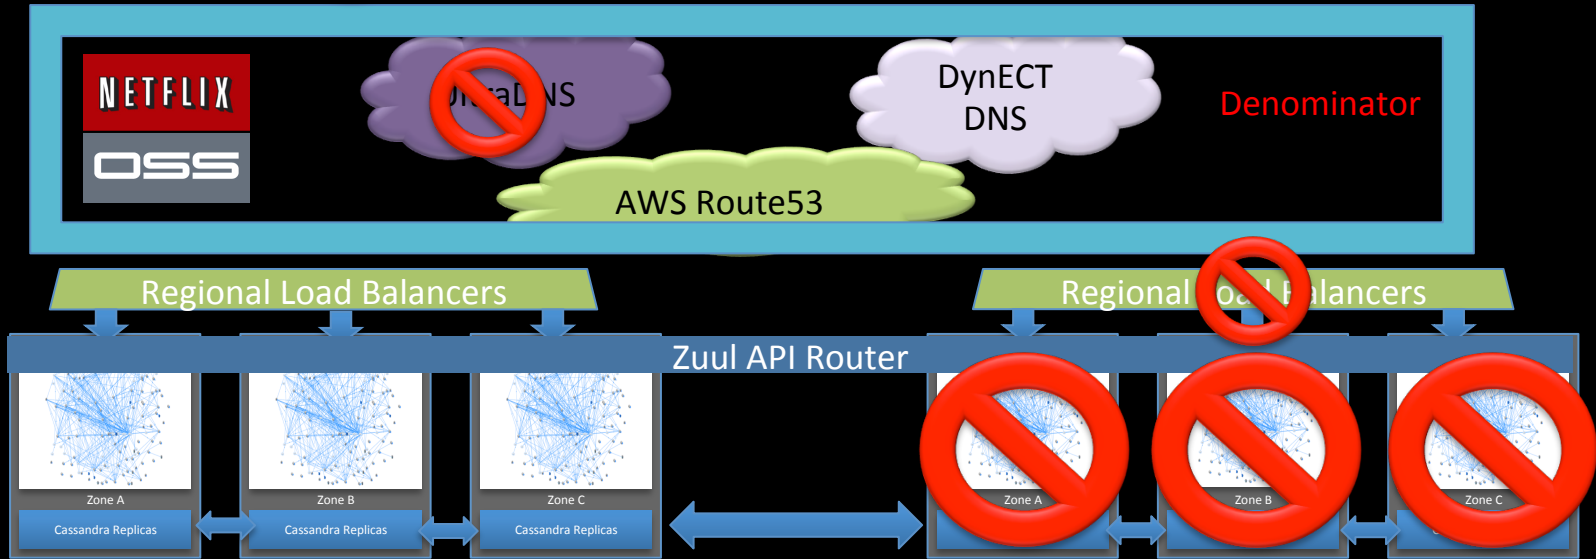
# Archaius library – configuration management



# Denominator: DNS for Multi-Region

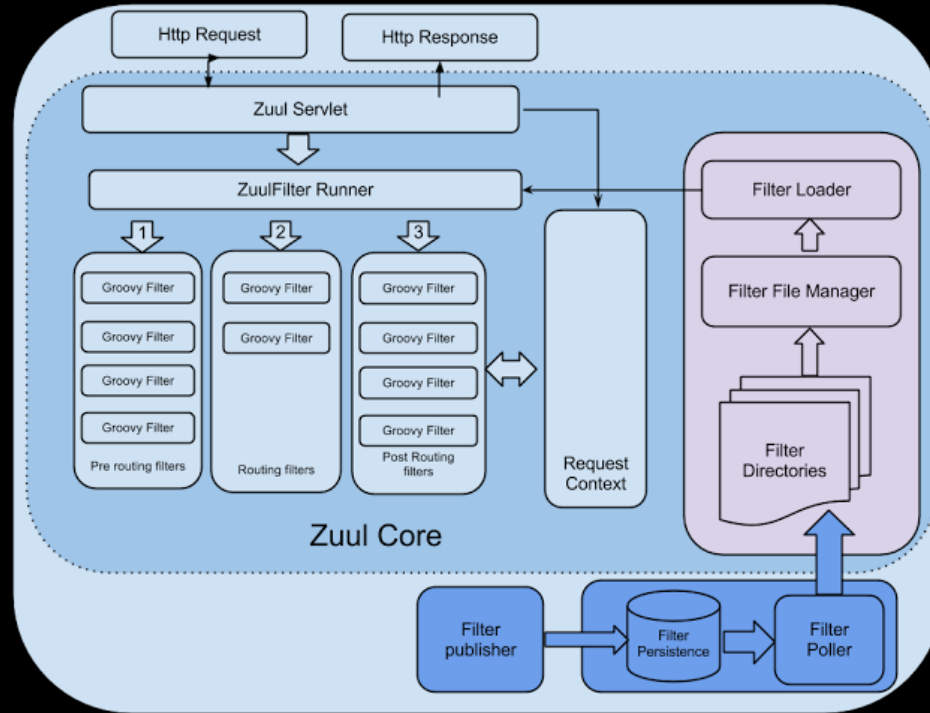


## Availability



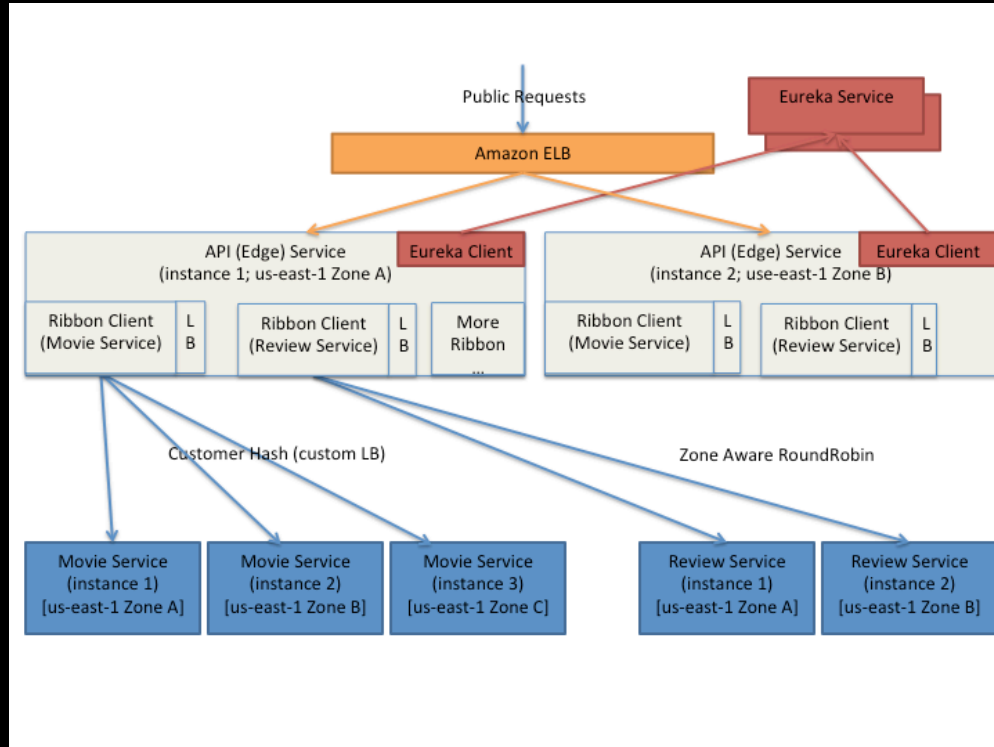
Denominator – manage traffic via multiple DNS providers with Java code

# Zuul – Smart and Scalable Routing Layer



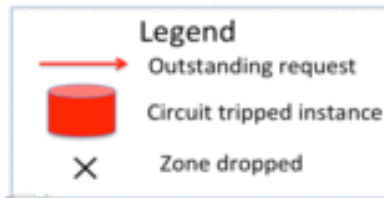


# Ribbon library for internal request routing



# Ribbon – Zone Aware LB

Average Active Requests = outstanding requests / (total instances - tripping instances)



# Karyon - Common server container



- Bootstrapping
  - Dependency & Lifecycle management via Governor.
  - Service registry via Eureka.
  - Property management via Archaius
  - Hooks for Latency Monkey testing
  - Preconfigured status page and heathcheck servlets

# Karyon



- Embedded Status Page Console
  - Environment
  - Eureka
  - JMX

The screenshot displays the Karyon Embedded Status Page Console for the application 'PLATFORMSERVICE v1.0 us-east-1\_prod'. The interface includes a navigation bar with tabs for 'Astranax', 'Counters', 'Debug Data', 'Diagnostics', 'Discovery', 'Environment', 'Info', 'Jars', 'JMX', 'NWS', 'Properties', 'SLA', and 'Tracers'. A 'Refresh' button and a 'Machine Readable' filter are also present. The left sidebar shows a tree view of the application's structure, with 'PLATFORM' expanded to show 'PLATFORMSERVICE' and 'LBStats\_AvailableZones' selected. The main content area displays the JMX attributes for 'com.netflix.servo.name=LBStats\_AvailableZones,'. The attributes are listed in a table with columns for 'Key', 'Name', and 'Value'. The table shows six entries, all with a 'value' name and various values representing availability zones. The status bar at the bottom indicates 'Error: Visible: 898 Total: 898 Last updated: Mon Mar 04 2013 23:13:52'.

Key	Name	Value
{\"class\":\"LoadBalancerStats\",\"id\":\"akmsclient\",\"level\":\"INFO\",\"type\":\"INFORMATIONAL\"}	value	[us-east-1e, us-east-1d, us-east-1c]
{\"class\":\"LoadBalancerStats\",\"id\":\"atlas_publish\",\"level\":\"INFO\",\"type\":\"INFORMATIONAL\"}	value	[us-east-1e, us-east-1d, us-east-1c]
{\"class\":\"LoadBalancerStats\",\"id\":\"chukwatrackerclient\",\"level\":\"INFO\",\"type\":\"INFORMATIONAL\"}	value	[us-east-1e, us-east-1d, us-east-1c]
{\"class\":\"LoadBalancerStats\",\"id\":\"defaultRestClient\",\"level\":\"INFO\",\"type\":\"INFORMATIONAL\"}	value	[]
{\"class\":\"LoadBalancerStats\",\"id\":\"epic_publish\",\"level\":\"INFO\",\"type\":\"INFORMATIONAL\"}	value	[us-east-1e, us-east-1d, us-east-1c]
{\"class\":\"LoadBalancerStats\",\"id\":\"epicplugin\",\"level\":\"INFO\",\"type\":\"INFORMATIONAL\"}	value	[us-east-1e, us-east-1d, us-east-1c]

# Karyon

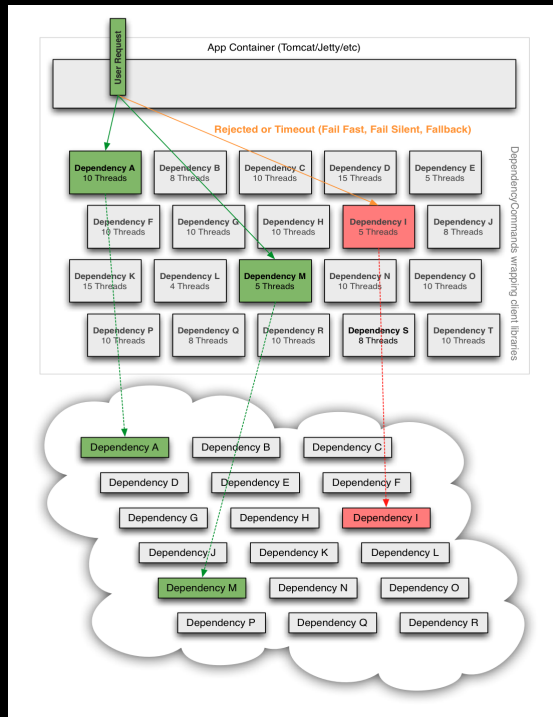


- Sample Service using Karyon available as "Hello-netflix-oss" on github

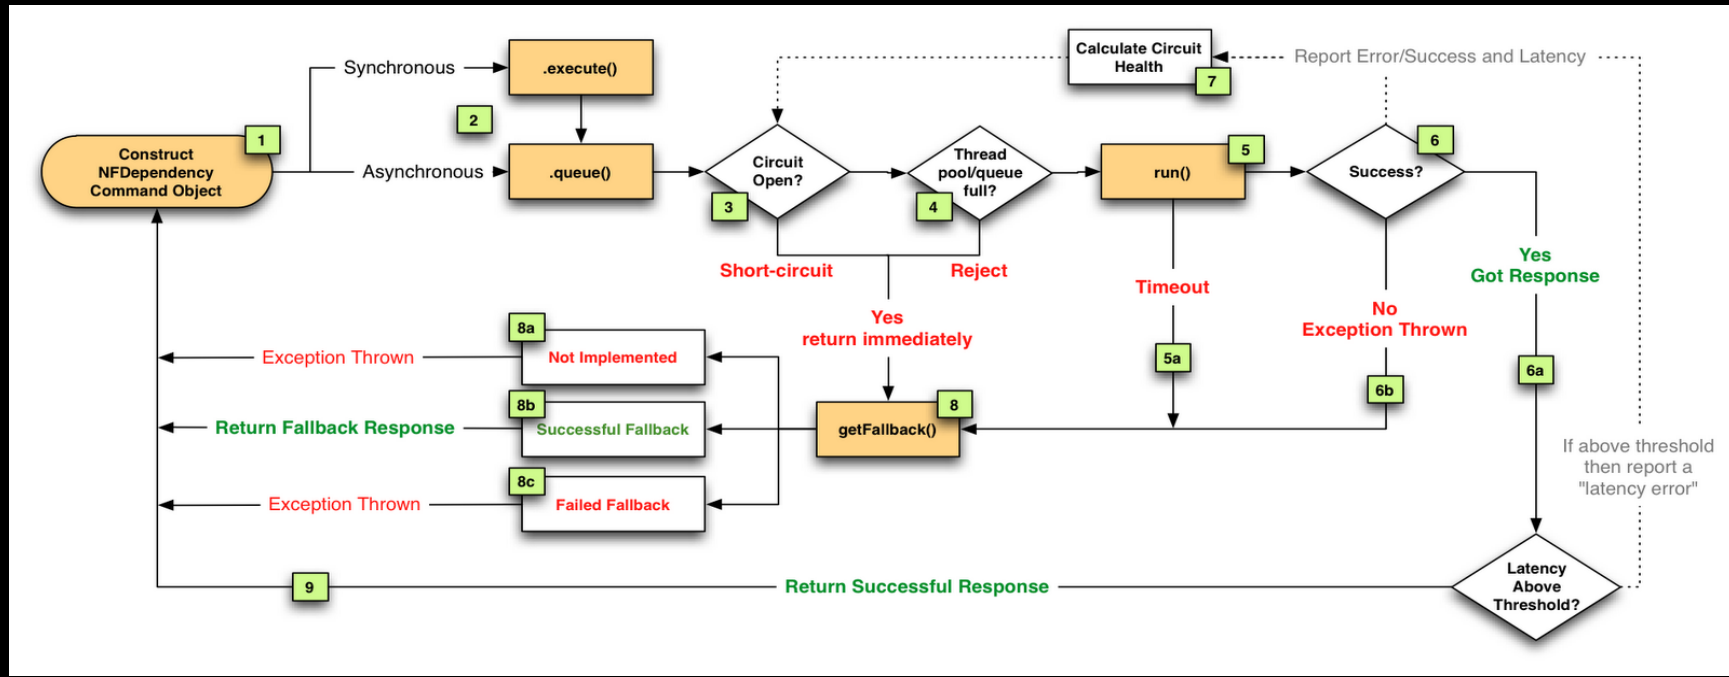
A screenshot of a web browser window. The address bar shows the URL "localhost:8989/hello-netflix-oss/rest/v1/hello/to/newbee". The browser's tab bar contains several tabs: "Netflix", "ReadMe", "XFINITY", "Utilities", "Ranch", "Tools", "Tahoe", "Wishlist", and "Add to Buyos". The main content area displays a JSON response: 

```
{  
  Message: "Hello newbie from Netflix OSS"  
}
```

# Hystrix Circuit Breaker: Fail Fast -> recover fast

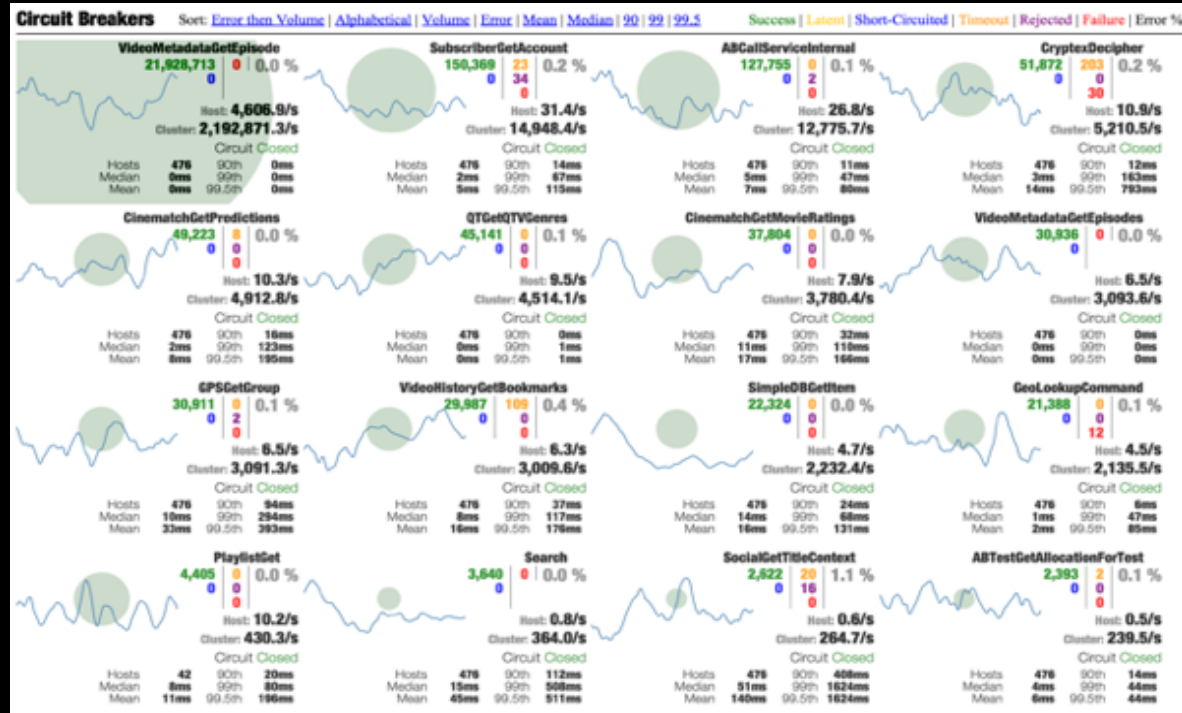


# Hystrix Circuit Breaker State Flow



# Turbine Dashboard

Per Second Update Circuit Breakers in a Web Browser





# Either you break it, or users will



NETFLIX

OSS

# Add some Chaos to your system



# Clean up your room! – Janitor Monkey

Works with Edda history to clean up after Asgard



# Conformity Monkey

Track and alert for old code versions and known issues  
Walks Karyon status pages found via Edda



# In use by many companies


The screenshot shows the Netflix OSS website header with the text "NETFLIX OSS Netflix Open Source Software Center". Below the header are navigation tabs for "Repositories", "Commit Timeline", and "Powered By NetflixOSS". A section titled "These companies are using and contributing to Netflix OSS Components" displays a grid of logos for various companies including KIXEYE, Yammer, FullContact, flipkart.com, globus genomics, Riot GAMES, coursera, yelp, Hotels.com, MORTAR, AnsWerS, YAHOO!, EUCALYPTUS, StumbleUpon, Maginatics, UserEvents, bazaarvoice, OpenSCG, and SUNCORP GROUP.

[netflixoss@netflix.com](mailto:netflixoss@netflix.com)

# http://www.meetup.com/Netflix-Open-Source-Platform/

## Netflix Open Source Platform

Home Members Sponsors Photos Pages Discussions More Gro



Change photo

**Los Gatos, CA**  
Founded Jan 2, 2013

About us...

Members	1,031
Group reviews	19
Upcoming	1

### Welcome!

CHANGE

+ SCHEDULE A NEW MEETUP

Upcoming 1 Past Calendar

#### Season 2 Episode 1

Netflix Inc  
100 Winchester Circle, Los Gatos, CA (map)

Wed Mar 12  
6:30 PM

I'M ATTENDING

192 attending  
58 spots left  
1 comment

Join us for the next installment of Netflix Open Source Meetups. Learn more about recently opened components, hear success stories from companies using NetflixOSS...

LEARN MORE

Hosted by: [Ruslan Meshenberg](#) (Organizer)

# Takeaways

*By embracing isolation and redundancy for availability  
we transformed our Cloud Infrastructure to be resilient against  
Region-wide outages*

*Use NetflixOSS to scale your Enterprise  
Contribute to existing github projects and add your own*

<http://netflix.github.com>

[@NetflixOSS](#)

[@rusmeshenberg](#)