

Globally Distributed Cloud Applications

Adrian Cockcroft
@adrianco

Netflix Inc.

Abstract

- Netflix on Cloud – What, Why and When
- Globally Distributed Availability Model
- Open Source Components
- Build Your Own PaaS

What Netflix Did

- Moved to SaaS
 - Corporate IT – OneLogin, Workday, Box, Evernote...
 - Tools – Pagerduty, AppDynamics, Elastic MapReduce
- Built our own PaaS
 - Customized to make our developers productive
 - When we started, we had little choice
- Moved incremental capacity to IaaS
 - No new datacenter space since 2008 as we grew
 - Moved our streaming apps to the cloud

Why Use Cloud?



Get stuck with wrong config

Wait Wait File tickets

Ask permission Wait Wait

Wait Things we don't do Wait

Run out of space/power

Plan capacity in advance

Have meetings with IT Wait

NETFLIX

Keeping up with Developer Trends

In production
at Netflix

- Big Data/Hadoop 2009
- AWS Cloud 2009
- Application Performance Management 2010
- Integrated DevOps Practices 2010
- Continuous Integration/Delivery 2010
- NoSQL 2010
- Platform as a Service; Fine grain SOA 2010
- Social coding, open development/github 2011

NETFLIX

AWS specific feature dependence....

Portability vs. Functionality

- Portability – the Operations focus
 - Avoid vendor lock-in
 - Support datacenter based use cases
 - Possible operations cost savings
- Functionality – the Developer focus
 - Less complex test and debug, one mature supplier
 - Faster time to market for your products
 - Possible developer cost savings

Functional PaaS

- IaaS base - all the features of AWS
 - Very large scale, mature, global, evolving rapidly
 - ELB, Autoscale, VPC, SQS, EIP, EMR, etc, etc.
 - Large files (TB) and multipart writes in S3
- Functional PaaS – Netflix added features
 - Very large scale, mature, flexible, customizable
 - Asgard console, Monkeys, Big data tools
 - Cassandra/Zookeeper data store automation



*Don't let the roadie
write the set list!*

*(yes you do need all
those guitars on
tour...)*

Developers
choose Functional
over Portable

NETFLIX

Freedom and Responsibility

- Developers leverage cloud to get freedom
 - Agility of a single organization, no silos
- But now developers are responsible
 - For compliance, performance, availability etc.

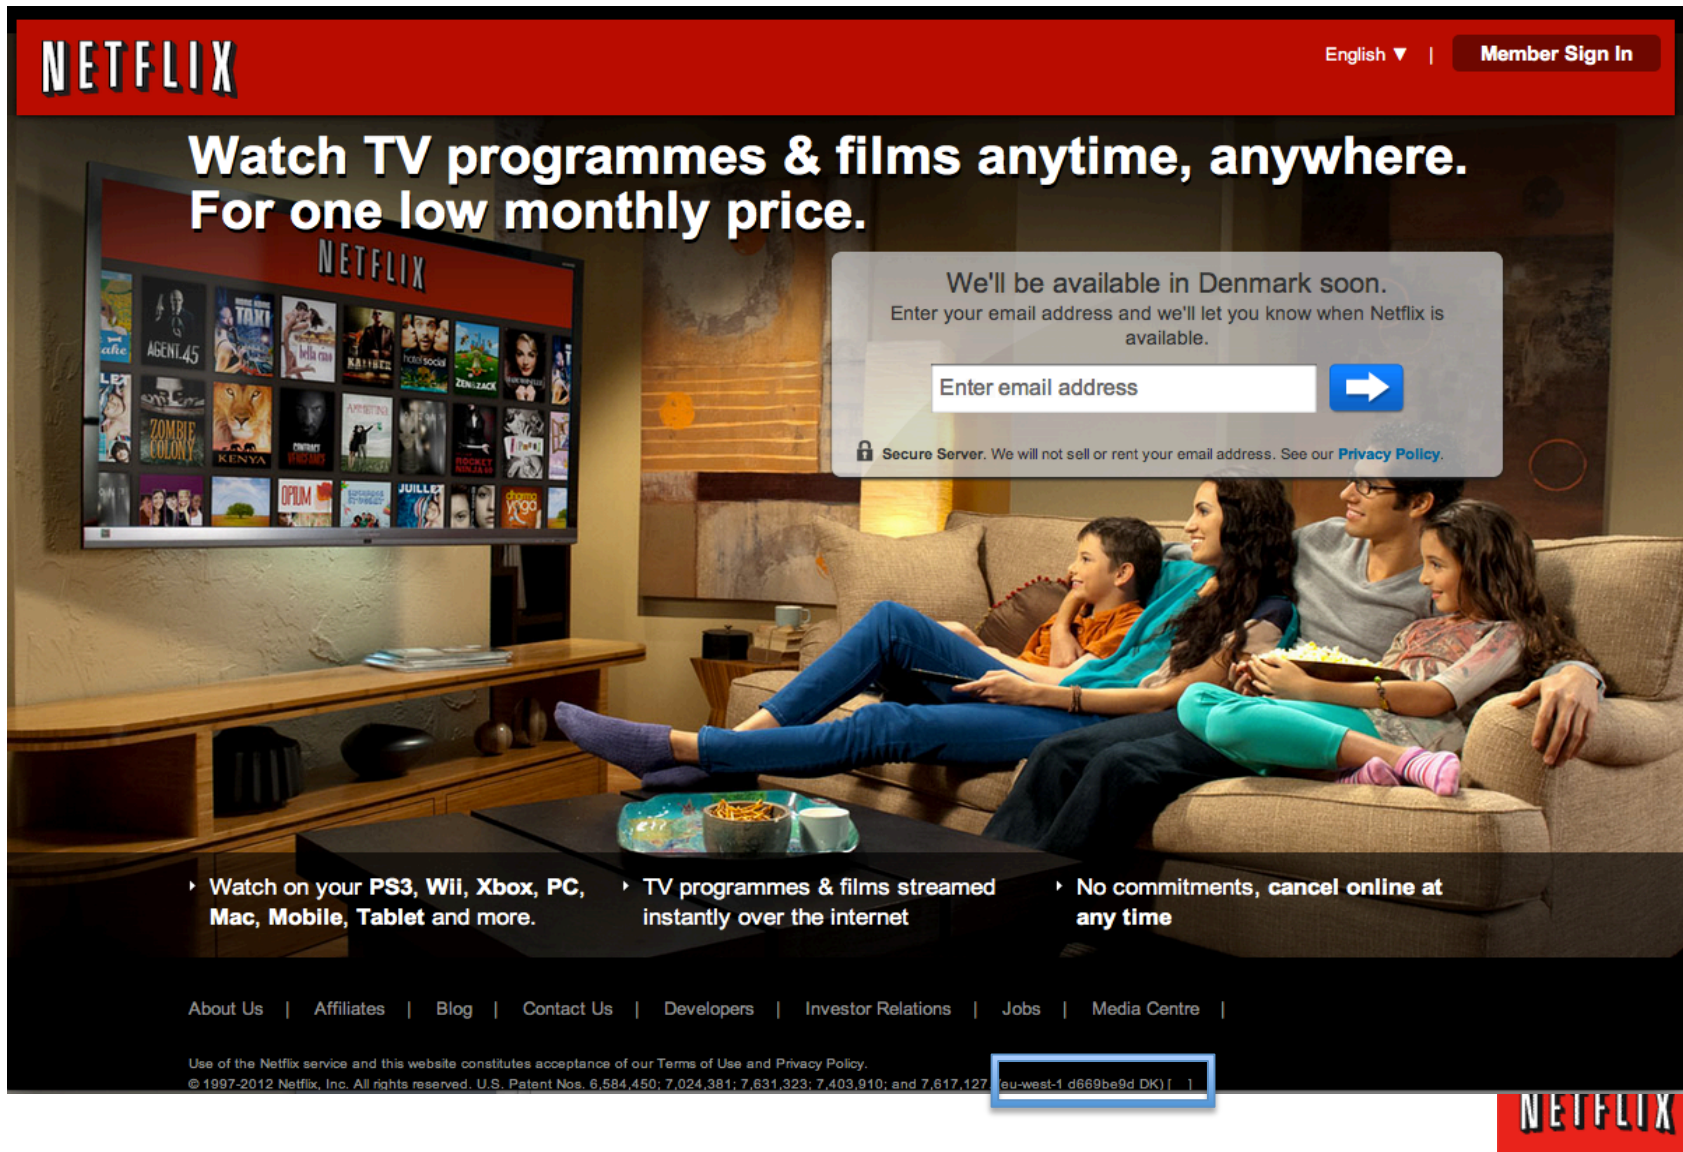
“As far as my rehab is concerned, it is within my ability to change and change for the better - Eddie Van Halen”

The Netflix Streaming Service

Coming soon to Denmark (and
Norway, Sweden and Finland)



www.netflix.com in DK



The image is a screenshot of the Netflix website's Danish landing page. At the top, the Netflix logo is on the left, and 'English' with a dropdown arrow and a 'Member Sign In' button are on the right. The main headline reads 'Watch TV programmes & films anytime, anywhere. For one low monthly price.' Below this, a large image shows a family of four sitting on a couch in a living room, watching a TV. The TV screen displays the Netflix interface with various movie and show thumbnails. Overlaid on the right side of the image is a grey sign-up box. It contains the text 'We'll be available in Denmark soon. Enter your email address and we'll let you know when Netflix is available.' Below this text is a white input field labeled 'Enter email address' and a blue button with a white right-pointing arrow. At the bottom of the sign-up box is a small lock icon followed by the text 'Secure Server. We will not sell or rent your email address. See our [Privacy Policy](#).' Below the main image, there are three bullet points: 'Watch on your PS3, Wii, Xbox, PC, Mac, Mobile, Tablet and more.', 'TV programmes & films streamed instantly over the internet', and 'No commitments, cancel online at any time'. The footer contains a navigation menu with links: 'About Us', 'Affiliates', 'Blog', 'Contact Us', 'Developers', 'Investor Relations', 'Jobs', and 'Media Centre'. Below the navigation menu is a small line of text: 'Use of the Netflix service and this website constitutes acceptance of our Terms of Use and Privacy Policy. © 1997-2012 Netflix, Inc. All rights reserved. U.S. Patent Nos. 6,584,450; 7,024,381; 7,631,323; 7,403,910; and 7,617,127'. To the right of this text is a blue-bordered box containing the text 'eu-west-1 d669be9d DK) []'. The Netflix logo is also present in the bottom right corner.

NETFLIX

English ▾ | Member Sign In

Watch TV programmes & films anytime, anywhere.
For one low monthly price.

We'll be available in Denmark soon.
Enter your email address and we'll let you know when Netflix is available.

Enter email address

Secure Server. We will not sell or rent your email address. See our [Privacy Policy](#).

- ▶ Watch on your PS3, Wii, Xbox, PC, Mac, Mobile, Tablet and more.
- ▶ TV programmes & films streamed instantly over the internet
- ▶ No commitments, cancel online at any time

About Us | Affiliates | Blog | Contact Us | Developers | Investor Relations | Jobs | Media Centre |

Use of the Netflix service and this website constitutes acceptance of our Terms of Use and Privacy Policy.
© 1997-2012 Netflix, Inc. All rights reserved. U.S. Patent Nos. 6,584,450; 7,024,381; 7,631,323; 7,403,910; and 7,617,127 eu-west-1 d669be9d DK) []

NETFLIX

Member Web Site

NETFLIXAdrian Cockcroft ▾ | Your Account & Help

Watch Instantly

Just for Kids

Browse DVDs

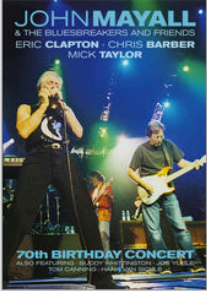
Your Queue

Taste Profile

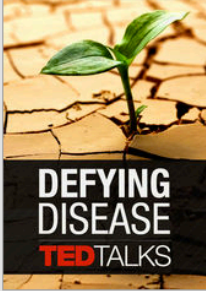
Genres ▾New ArrivalsInstantly to your TV

Recently Watched

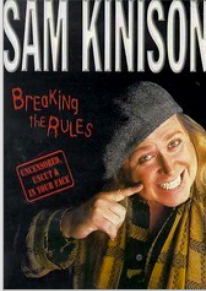
Top 10 for Adrian



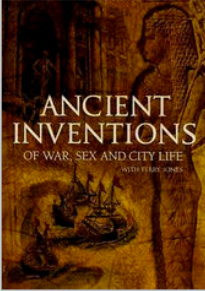
JOHN MAYALL
& THE BLUESBREAKERS AND FRIENDS
ERIC CLAPTON · CHRIS BARBER
MICK TAYLOR
70th BIRTHDAY CONCERT



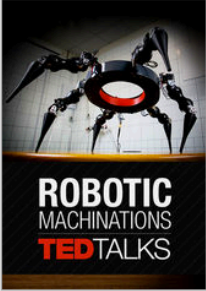
DEFYING
DISEASE
TED TALKS




SAM KINISON
Breaking
the RULES
CENSORED
MUTY &
LOVE PAT




ANCIENT
INVENTIONS
OF WAR, SEX AND CITY LIFE
WITH YVES ANJES




ROBOTIC
MACHINATIONS
TED TALKS







that Mitchell
and Webb look



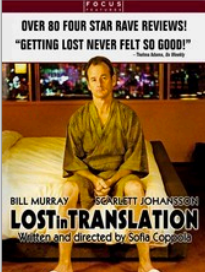
Bartleby

 Friends' Favorites

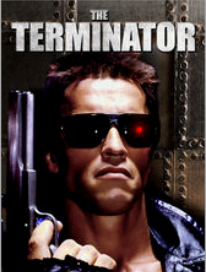
Based on these friends:




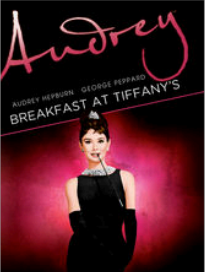
Breaking
Bad



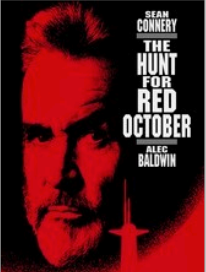
OVER 80 FOUR STAR RAVE REVIEWS!
"GETTING LOST NEVER FELT SO GOOD!"
BILL MURRAY · SCARLETT JOHANSSON
LOST in TRANSLATION
Written and directed by Sofia Coppola




THE
TERMINATOR



Audrey
BREAKFAST AT TIFFANY'S
AUDREY HEPBURN · GEORGE PEPPARD



SEAN CONNERY
THE HUNT FOR
RED OCTOBER
ALEC BALDWIN



ROBIN WILLIAMS
MATT DAMON
GOOD WILL HUNTING

NETFLIX

Content Delivery Service

Distributed storage nodes controlled by Netflix cloud services

NETFLIX

Open Connect

Overview

FAQ

Peering Information

> Hardware Design

Software Design

Deployment Guide

ISP Inquiry

Open Connect Appliance Hardware

Objectives

When designing the Open Connect Appliance Hardware, we focused on these fundamental design goals:

- Very high storage density without sacrificing space and power efficiency. Our target was fitting 100 terabytes into a 4u chassis that is less than 2' deep.
- High throughput: 10 Gbps throughput via an optical network connection.
- Very low field maintenance: the appliance must tolerate a variety of hardware failures including hard drives, network optics, and power supply units.
- Simple racking and installation. Front mounted power and network ports are the only things to connect at install time.



Open Connect Appliances are servers based on commodity PC components (similar to the model used by all large scale content delivery networks). We were influenced by the excellent write-ups from the [Backblaze](#) team, and use a custom chassis due to a lack of ready made options for a compact unit.

To achieve over 100 TB of storage, spinning hard drives provide the highest affordable density, in particular 36 3TB SATA units. The hard drives are not hot swappable, as we wish to avoid the operational burden of field service. For lower power utilization and simpler sourcing we select commodity units from two vendors and use software to manage failure modes and avoid field replacement. Dead drives reduce the total storage available for the system, but don't take it offline. We also add 1 TB of flash storage (2 solid state drives) for system files, logs and popular content. To augment the motherboard attached controller, we use two 16 port LSI SAS controller cards that connect directly to the SATA drives. This avoids I/O bottlenecks of SATA multipliers or SAS expanders, and also reduces system complexity.

From a compute point of view, the system has modest requirements moving bits from the storage to network packets on the interface. To reduce the power usage and hence also cooling requirement (which in turn reduces vibration from case fans) we use a single low power 4 core Intel Sandy Bridge CPU on a small form factor [Supermicro](#) mATX board with the full 32 GB of RAM installed.

We use redundant, hot swappable power supply units that have interchangeable AC and DC options for maximum installation flexibility. [Zippy](#) reversed the fan rotation of the units to allow mounting at the front of the case, and thus allow network and power connects to be positioned here.


The network card has two 10 Gbps modules, which can power a variety of SR and LR optic modules, for installation flexibility and scalable interconnection.

The following system was developed and first deployed at the end of 2011.

NETFLIX

Streaming Device API

KOKU Netflix Player by Roku




"My guess is that eventually, the streaming feature will be part of Blu-ray players and TVs. But for now, the Netflix Player by Roku strikes me as a great value for early adopters and film addicts alike."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Roku
Model	HD (N1000)
Variants	None
Main Color	Black
Product Type	Set-top box
Market (a)	US
Version	1.0
Date	May 20, 2008

MICROSOFT Xbox 360




"It's high-def Netflix streaming, and it's coming first to Microsoft's Xbox 360. Yes friends, when the all-new dashboard hits on November 19th, with it will come HD Netflix streaming for Xbox Live Gold members."

Engadget, Darren Murph
October 28, 2008

Cost Brand	Microsoft
Model	Xbox 360
Variants	None
Main Color	3.0GHz PPC Tri-Core Xenon
Date	Nov. 18, 2008

XBOX 360




"The most popular choice of the decade in the first survey... it's the only console you can't live without."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Microsoft
Model	Xbox 360
Variants	None
Main Color	3.0GHz PPC Tri-Core Xenon
Date	Nov. 18, 2008

LG BD390




"Streaming Netflix is one of the main selling points of the LG BD390."

LG BD390 for sale
April, 2009

Cost Brand	LG Electronics
Model	BD390
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Apr. 12, 2009

Samsung HT-BD1200




"The 2008 Samsung HT-BD1200 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Samsung
Model	HT-BD1200
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	May 20, 2008

LG LHB953




"If the internet is the key to the future of the home theater, then the LG LHB953 is the key to the future of the home theater. It's a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	LG
Model	LHB953
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	May 20, 2008

LG LHB953




"If the internet is the key to the future of the home theater, then the LG LHB953 is the key to the future of the home theater. It's a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	LG
Model	LHB953
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	May 20, 2008

Samsung HT-BD2820




"The 2008 Samsung HT-BD2820 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Samsung
Model	HT-BD2820
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	May 20, 2008

LG BD390




"The 2008 LG BD390 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	LG
Model	BD390
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	May 20, 2008

INSIGNIA NS-BRDVD3




"The 2008 Insignia NS-BRDVD3 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	INSIGNIA
Model	NS-BRDVD3
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	May 20, 2008

Sony PlayStation 3




"The 2008 Sony PlayStation 3 console is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Sony
Model	PS3
Variants	None
Main Color	Black
Product Type	Game Console
Market (a)	US
Version	1.0
Date	Nov. 6, 2009

SONY KDL-NBR9




"The 2008 Sony KDL-NBR9 TV is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	SONY
Model	KDL-NBR9
Variants	None
Main Color	Black
Product Type	TV
Market (a)	US
Version	1.0
Date	Nov. 6, 2009

SONY KDL-NBR9




"The 2008 Sony KDL-NBR9 TV is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	SONY
Model	KDL-NBR9
Variants	None
Main Color	Black
Product Type	TV
Market (a)	US
Version	1.0
Date	Nov. 6, 2009

SONY BDP-N460




"The 2008 Sony BDP-N460 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	SONY
Model	BDP-N460
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Nov. 6, 2009

INSIGNIA NS-BRDVD3




"The 2008 Insignia NS-BRDVD3 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	INSIGNIA
Model	NS-BRDVD3
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Nov. 6, 2009

Nintendo Wii




"The 2008 Nintendo Wii console is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Nintendo
Model	Wii
Variants	None
Main Color	Black
Product Type	Game Console
Market (a)	US
Version	1.0
Date	Mar. 25, 2010

VIZIO SV422XVT




"The 2008 Vizio SV422XVT TV is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Vizio
Model	SV422XVT
Variants	None
Main Color	Black
Product Type	TV
Market (a)	US
Version	1.0
Date	Jan. 2010

VIZIO VBR-200W




"The 2008 Vizio VBR-200W Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Vizio
Model	VBR-200W
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Jan. 2010

Apple iPad




"The 2008 Apple iPad tablet is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Apple
Model	iPad
Variants	None
Main Color	Black
Product Type	Tablet
Market (a)	US
Version	1.0
Date	Mar. 25, 2010

Samsung LED C8000




"The 2008 Samsung LED C8000 TV is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Samsung
Model	LED C8000
Variants	None
Main Color	Black
Product Type	TV
Market (a)	US
Version	1.0
Date	Mar. 1, 2010

PHILIPS BDP5110




"The 2008 Philips BDP5110 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	PHILIPS
Model	BDP5110
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Mar. 25, 2010

Samsung HT-C6900




"The 2008 Samsung HT-C6900 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Samsung
Model	HT-C6900
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Mar. 25, 2010

Samsung HT-C6900W




"The 2008 Samsung HT-C6900W Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	Samsung
Model	HT-C6900W
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Mar. 25, 2010

LG LK5500



"The 2008 LG LK5500 Blu-ray player is a great choice for those who want a high-end Blu-ray player with a built-in HD media player and also a great choice for those who want a high-end Blu-ray player with a built-in HD media player."

PCMag.com, Tim Gilman
May 20, 2008

Cost Brand	LG
Model	LK5500
Variants	None
Main Color	Black
Product Type	Blu-ray player
Market (a)	US
Version	1.0
Date	Mar. 25, 2010

Netfix Ready Devices

From: May 2008
To: May 2010

instant streaming ready


NETFLIX

Major Front End Services

- Non-member Web Site www.netflix.com
 - Marketing driven, sign up flow, SOX/PCI scope
- Member Web Site movies.netflix.com
 - Personalization driven
- CDNs for delivering bulk video/audio
 - Netflix CDN: openconnect.netflix.com
- API for external and device user interfaces
 - Mostly private APIs, limited public API docs at developer.netflix.com
- API for controlling video playback
 - DRM, QoS management, Bookmarks

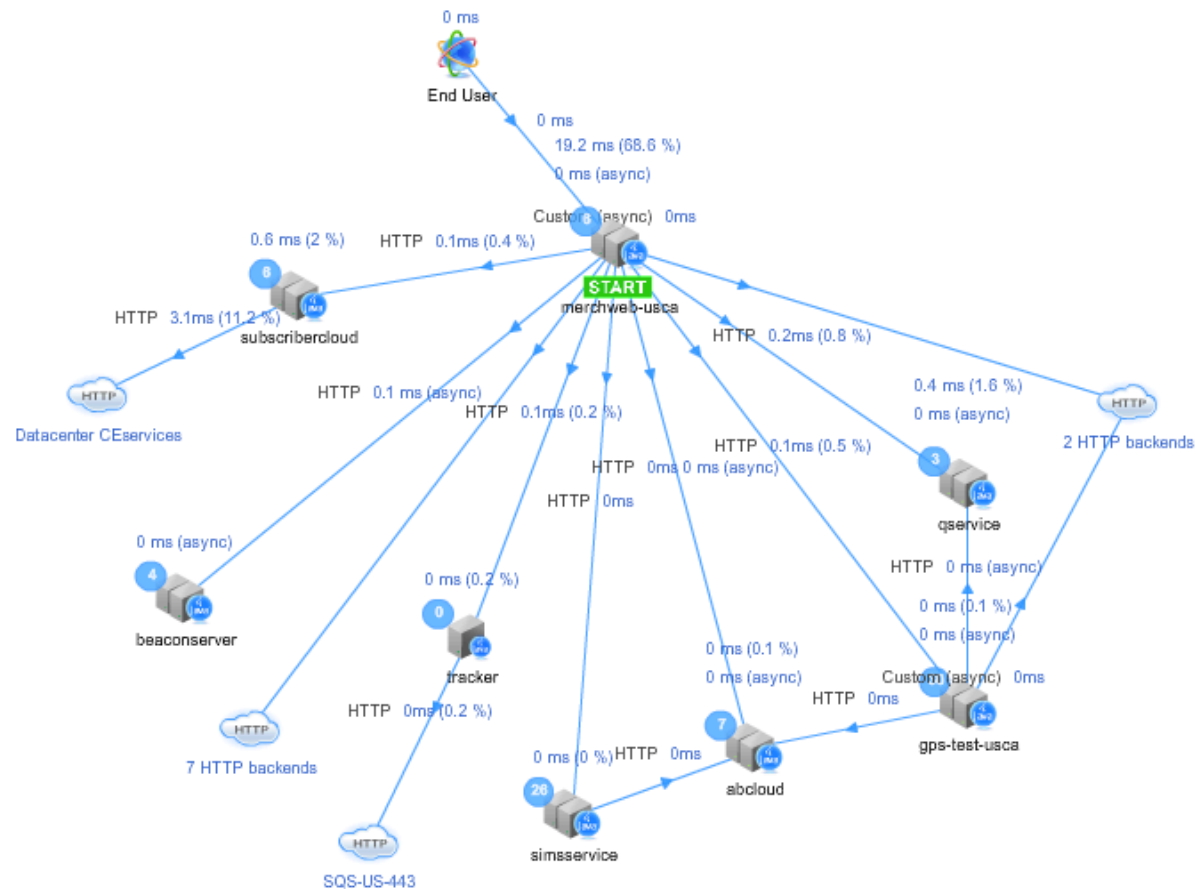
Component Services

(Simplified view using AppDynamics)



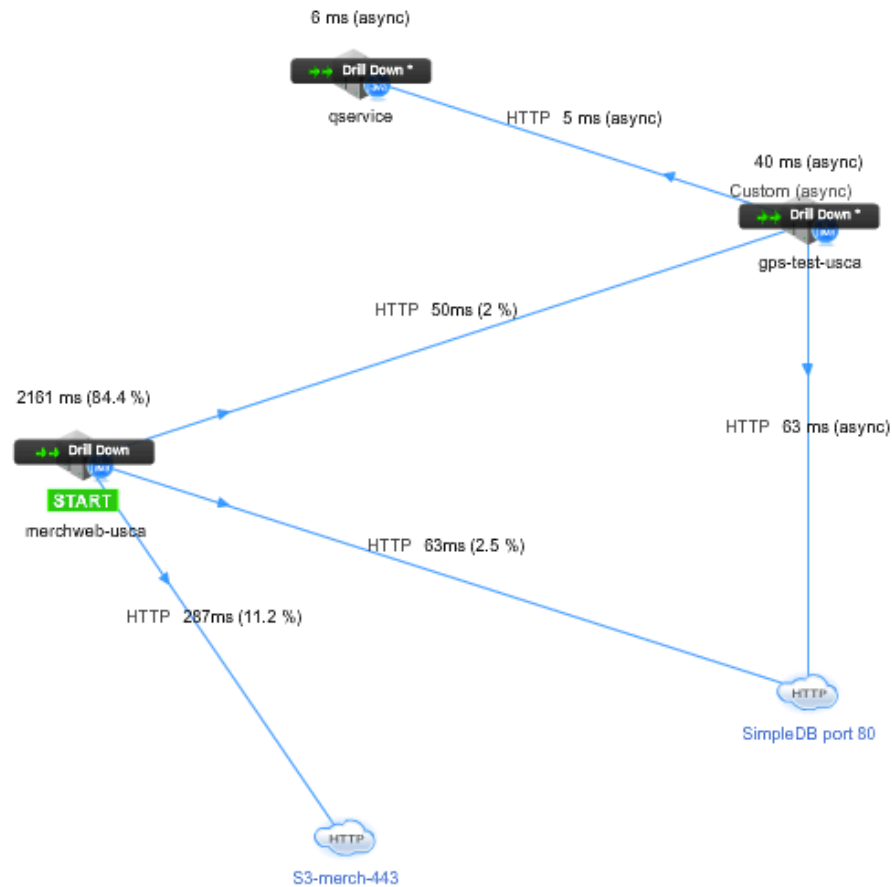
Web Server Dependencies Flow

(zooming in to one business transaction in test account)



One Request Snapshot

(captured because it was unusually slow)




Current Architectural Patterns for Availability

- Isolated Services
 - Resilient Business logic
- Three Balanced Availability Zones
 - Resilient to Infrastructure outage
- Triple Replicated Persistence
 - Durable distributed Storage
- Isolated Regions
 - US and EU don't take each other down

Isolated Services

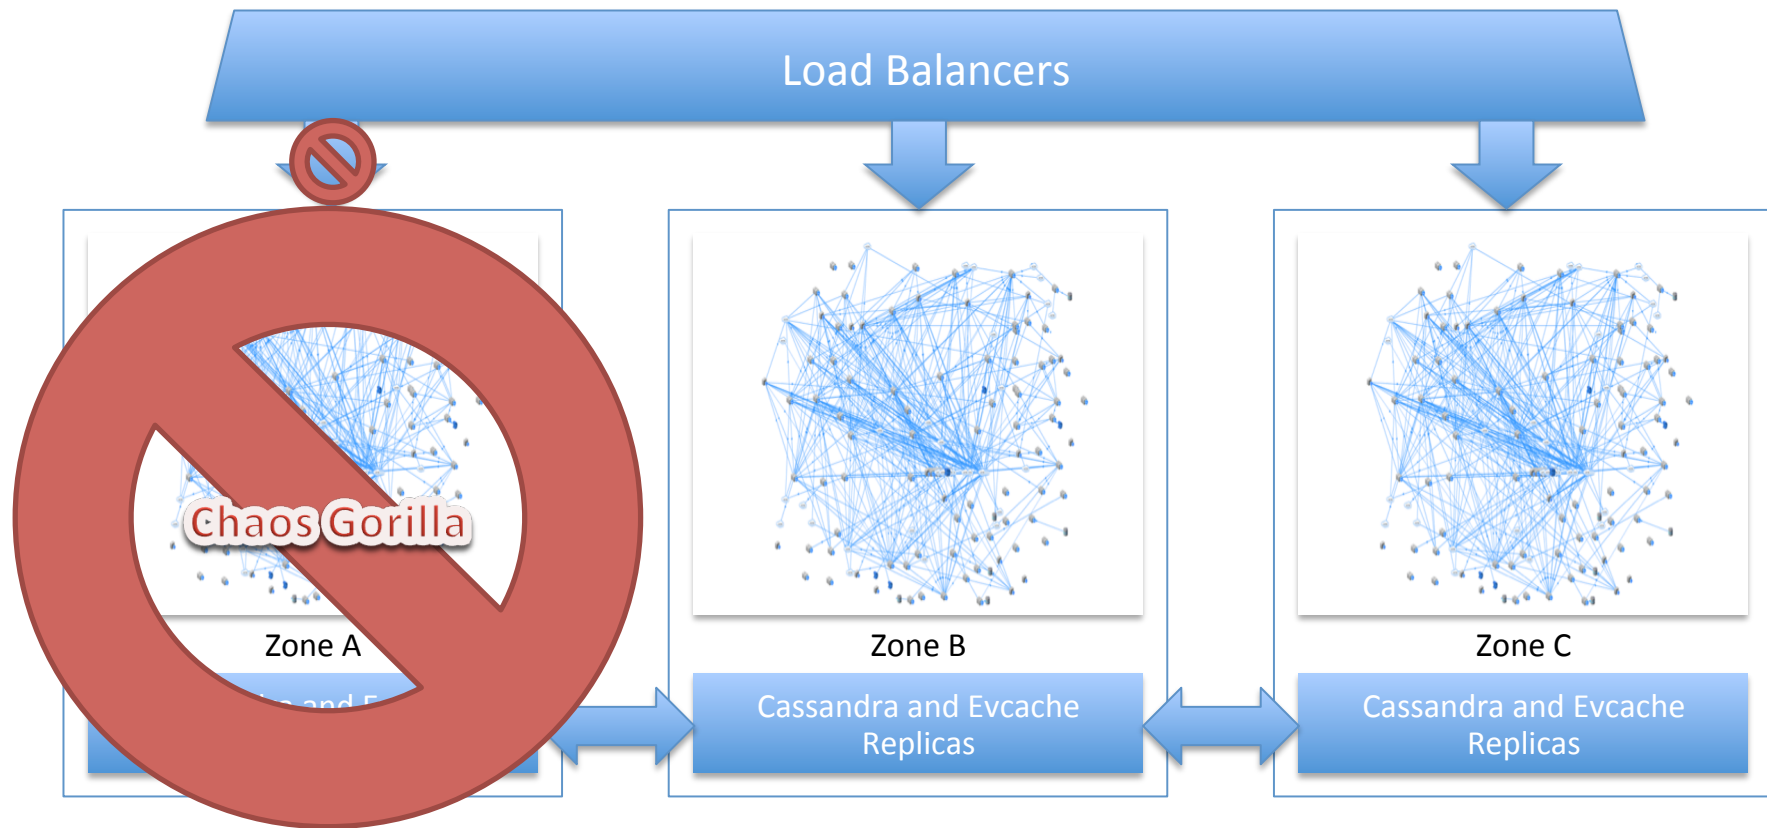
Test With Chaos Monkey, Latency Monkey



NETFLIX

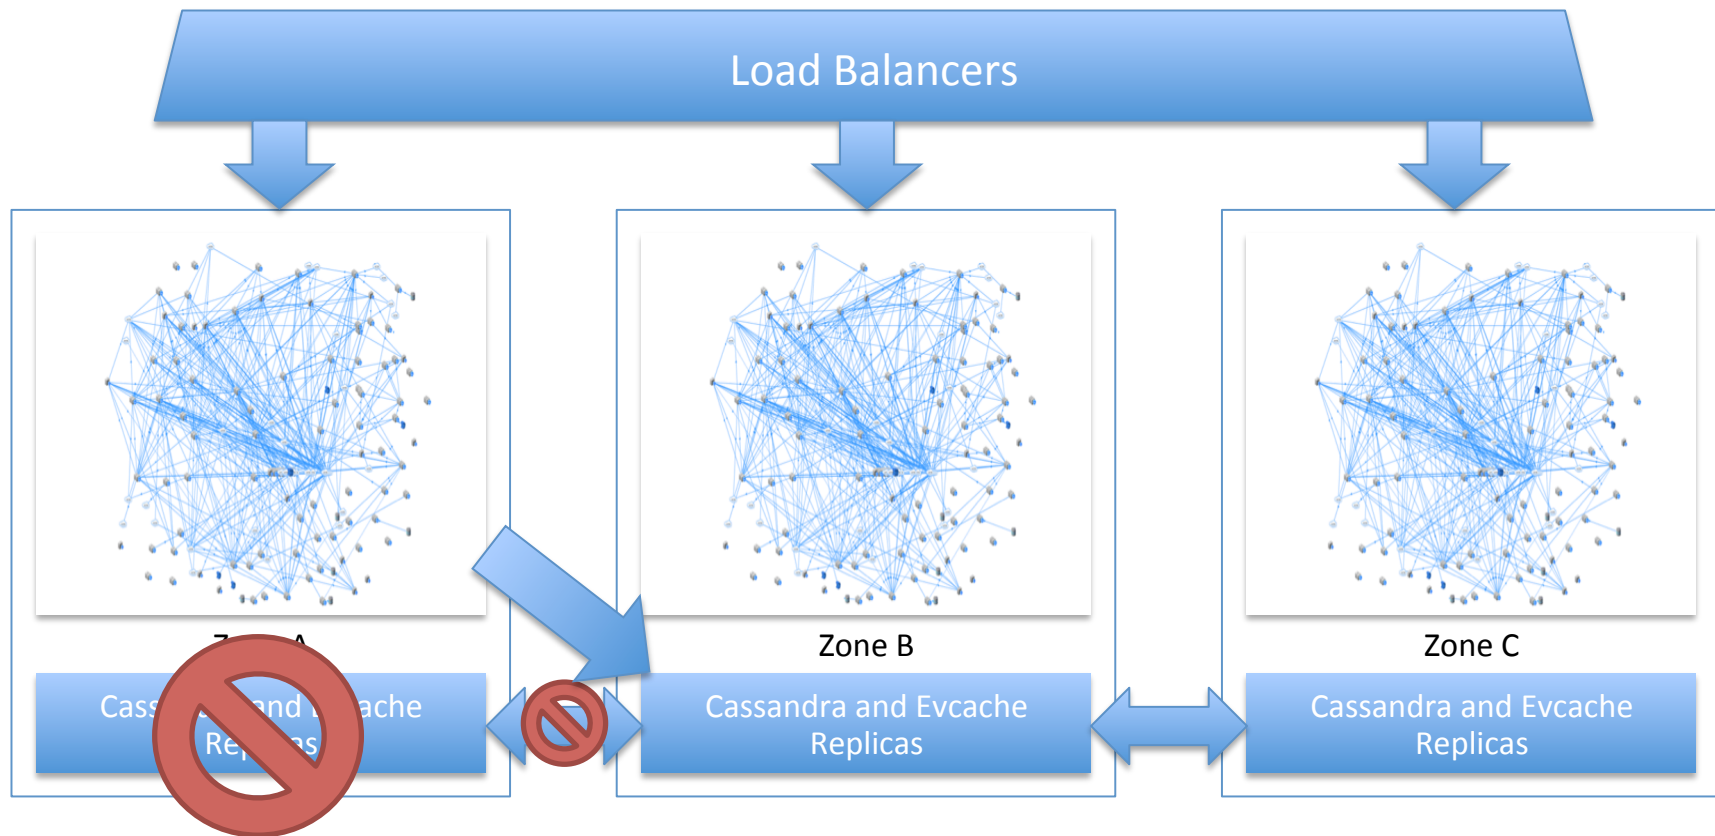
Three Balanced Availability Zones

Test with Chaos Gorilla

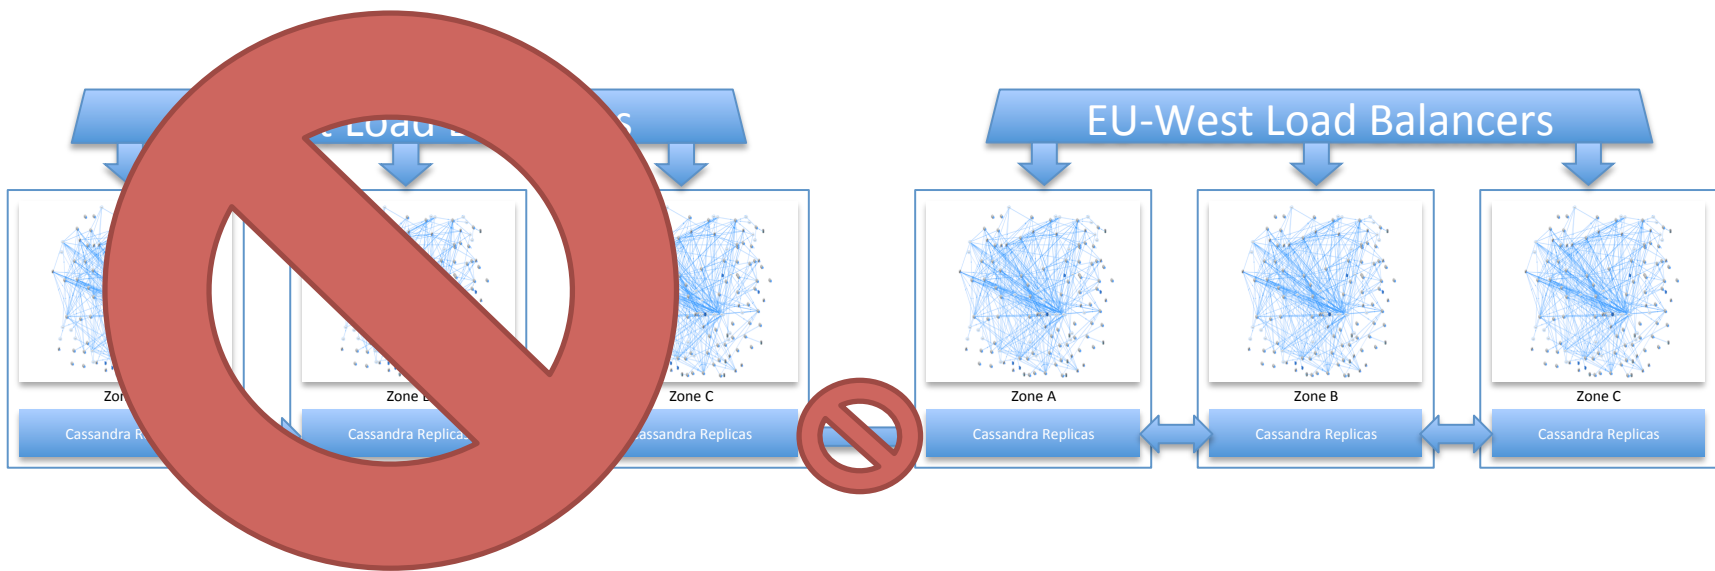


Triple Replicated Persistence

Cassandra maintenance drops individual replicas



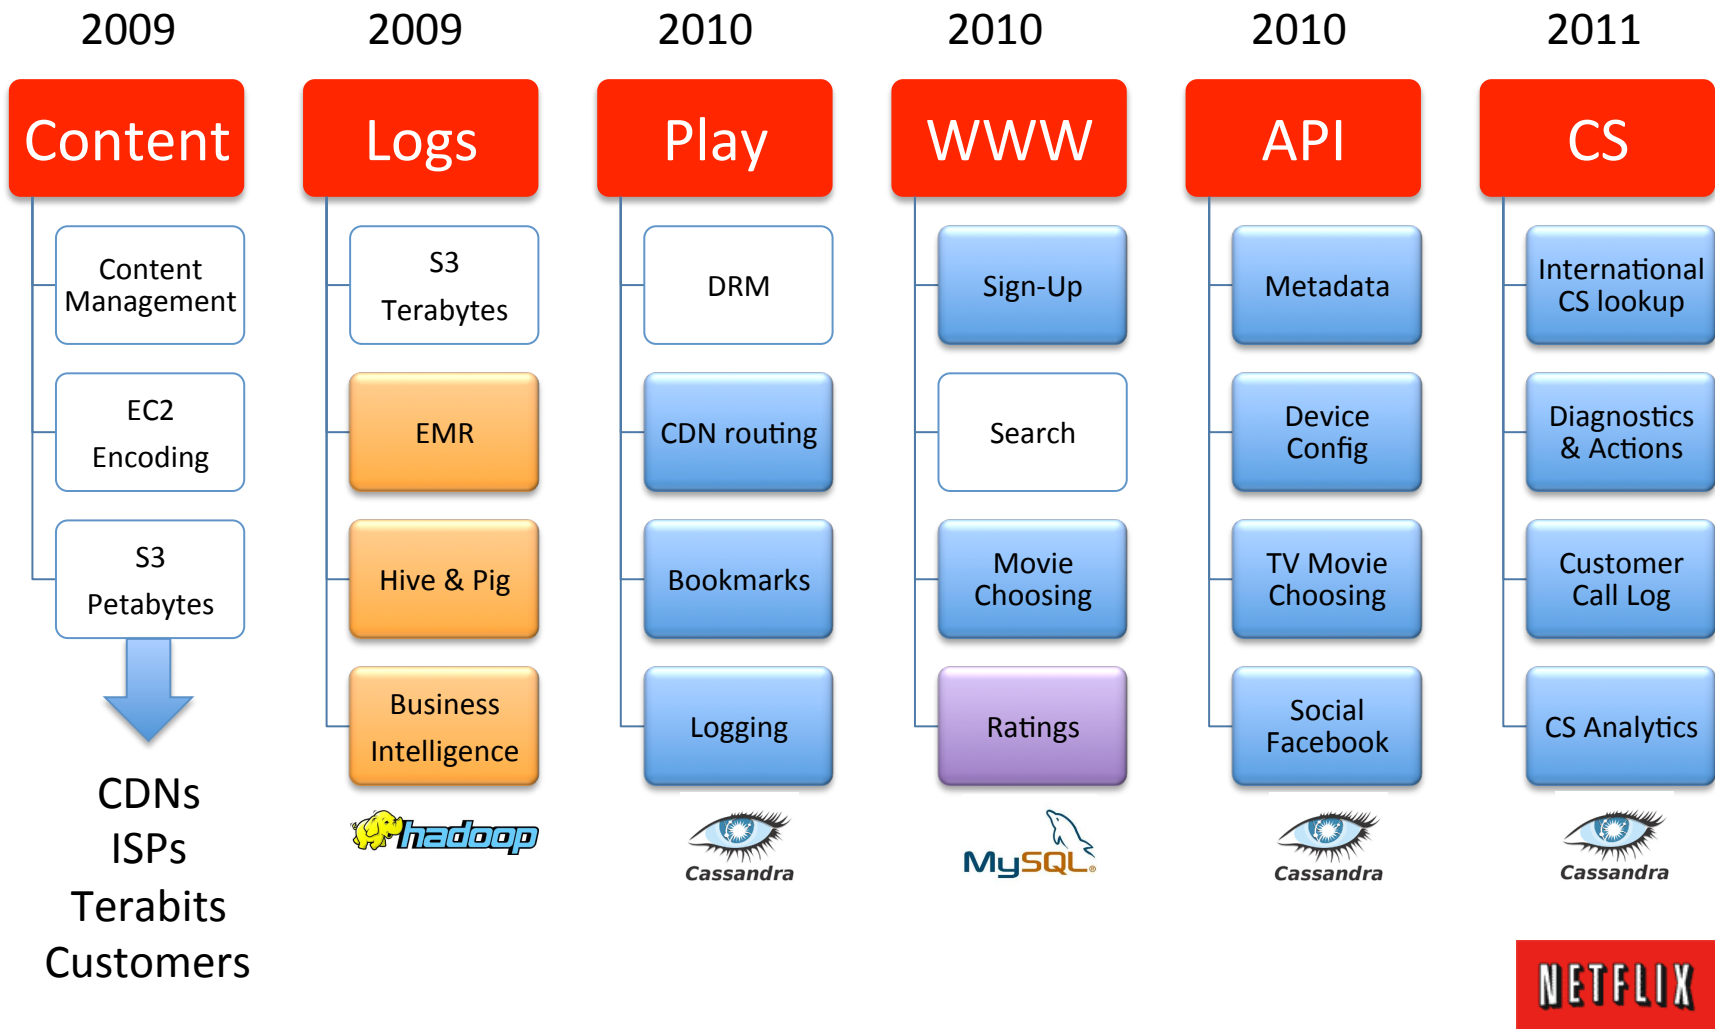
Isolated Regions



Failure Modes and Effects

Failure Mode	Probability	Mitigation Plan
Application Failure	High	Automatic degraded response
AWS Region Failure	Low	Wait for region to recover
AWS Zone Failure	Medium	Continue to run on 2 out of 3 zones
Datacenter Failure	Medium	Migrate more functions to cloud
Data store failure	Low	Restore from S3 backups
S3 failure	Low	Restore from remote archive

Netflix Deployed on AWS



Cloud Architecture Patterns

Where do we start?



Datacenter to Cloud Transition Goals

- **Faster**
 - **Lower latency** than the equivalent datacenter web pages and API calls
 - Measured as mean and 99th percentile
 - For both first hit (e.g. home page) and in-session hits for the same user
- **Scalable**
 - **Avoid needing any more datacenter capacity** as subscriber count increases
 - No central vertically scaled databases
 - Leverage AWS elastic capacity effectively
- **Available**
 - Substantially **higher robustness and availability** than datacenter services
 - Leverage multiple AWS availability zones
 - No scheduled down time, no central database schema to change
- **Productive**
 - Optimize **agility** of a large development team with automation and tools
 - Leave behind complex tangled datacenter code base (~8 year old architecture)
 - Enforce clean layered interfaces and re-usable components

Netflix Datacenter vs. Cloud Arch

Anti-Architecture

Central SQL Database

Distributed Key/Value NoSQL

Sticky In-Memory Session

Shared Memcached Session

Chatty Protocols

Latency Tolerant Protocols

Tangled Service Interfaces

Layered Service Interfaces

Instrumented Code

Instrumented Service Patterns

Fat Complex Objects

Lightweight Serializable Objects

Components as Jar Files

Components as Services

NETFLIX

Availability and Resilience

Chaos Monkey

<http://techblog.netflix.com/2012/07/chaos-monkey-released-into-wild.html>

- Computers (Datacenter or AWS) randomly die
 - Fact of life, but too infrequent to test resiliency
- Test to make sure systems are resilient
 - Allow any instance to fail without customer impact
- Chaos Monkey hours
 - Monday-Friday 9am-3pm random instance kill
- Application configuration option
 - Apps now have to opt-out from Chaos Monkey



Responsibility and Experience

- Make developers responsible for failures
 - Then they learn and write code that doesn't fail
- Use Incident Reviews to find gaps to fix
 - Make sure its not about finding “who to blame”
- Keep timeouts short, fail fast
 - Don't let cascading timeouts stack up
- Make configuration options dynamic
 - You don't want to push code to tweak an option

<http://techblog.netflix.com/2012/02/fault-tolerance-in-high-volume.html>



Build Your Own PaaS



Components

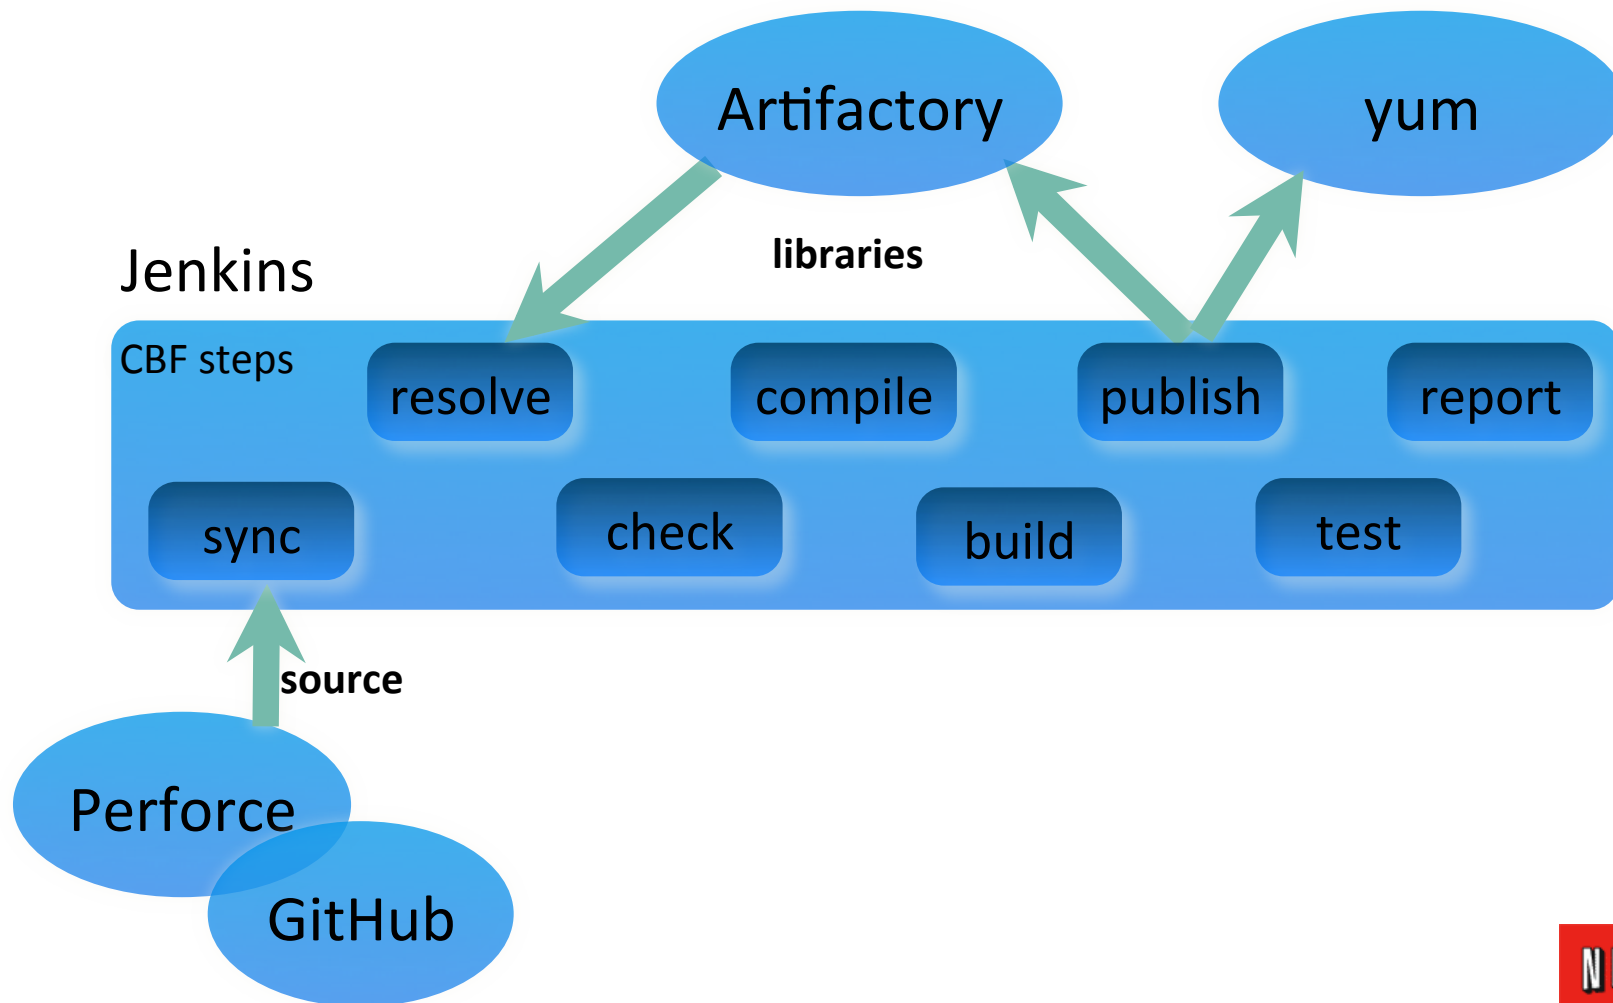
- Continuous build framework turns code into AMIs
- AWS accounts for test, production, etc.
- Cloud access gateway
- Service registry
- Configuration properties service
- Persistence services
- Monitoring, alert forwarding
- Backups, archives

Common Build Framework

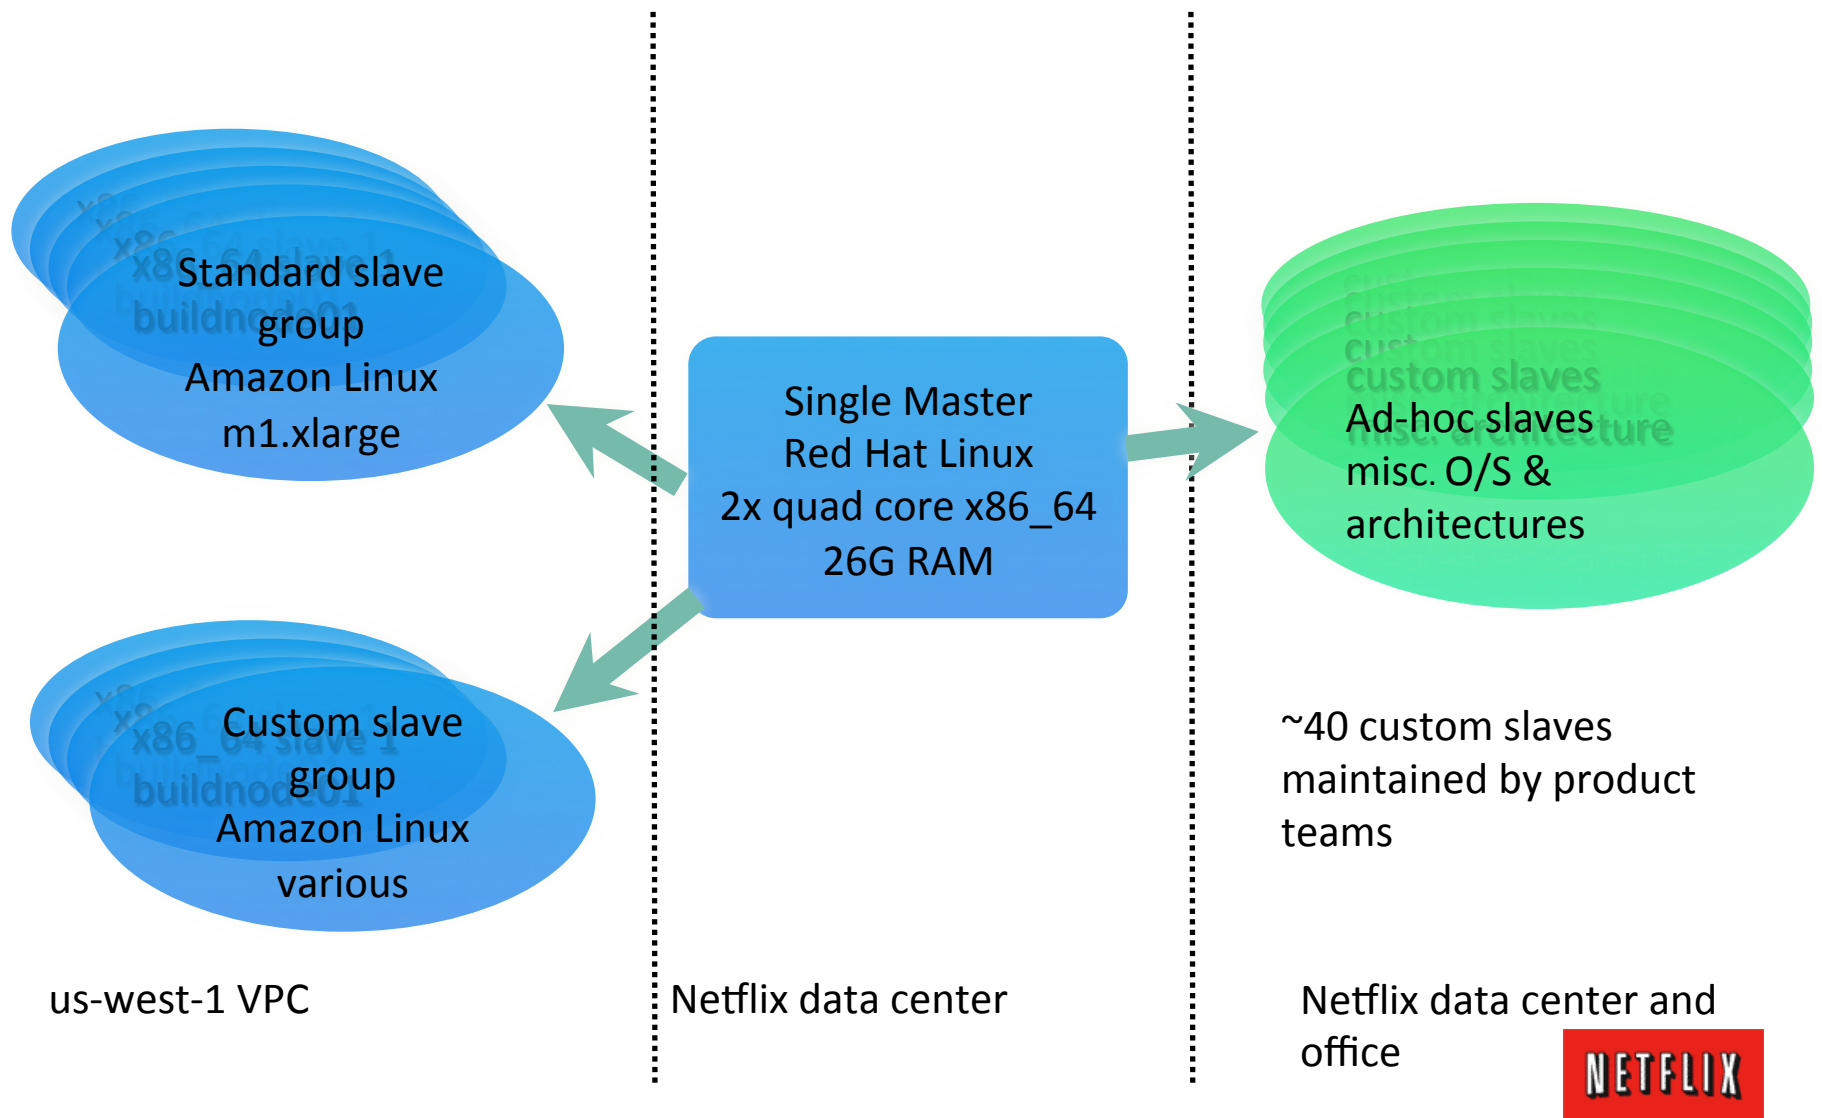
Extracted from
“Building and Deploying Netflix in the Cloud”
by @bmoyles and @garethbowles
On slideshare.net/netflix



Build Pipeline



Jenkins Architecture



Netflix Extensions to Jenkins

Job DSL plugin: allow jobs to be set up with minimal definition, using templates and a Groovy-based DSL (Gradle)

Housekeeping and maintenance processes implemented as Jenkins jobs, system Groovy scripts

The Bakery

- Create base AMIs
 - We have CentOS, Ubuntu and Windows base AMIs
 - All the generic code, apache, tomcat etc.
 - Standard system and application monitoring tools
 - Update ~monthly with patches and new versions
- Add yummy topping and bake
 - Jenkins builds app specific AMI including all code etc.
 - Bakery mounts EBS snapshot, installs and bakes
 - One bakery per region, delivers into paastest
 - Tweak config and publish AMI to paasprod

AWS Accounts



Accounts Isolate Concerns

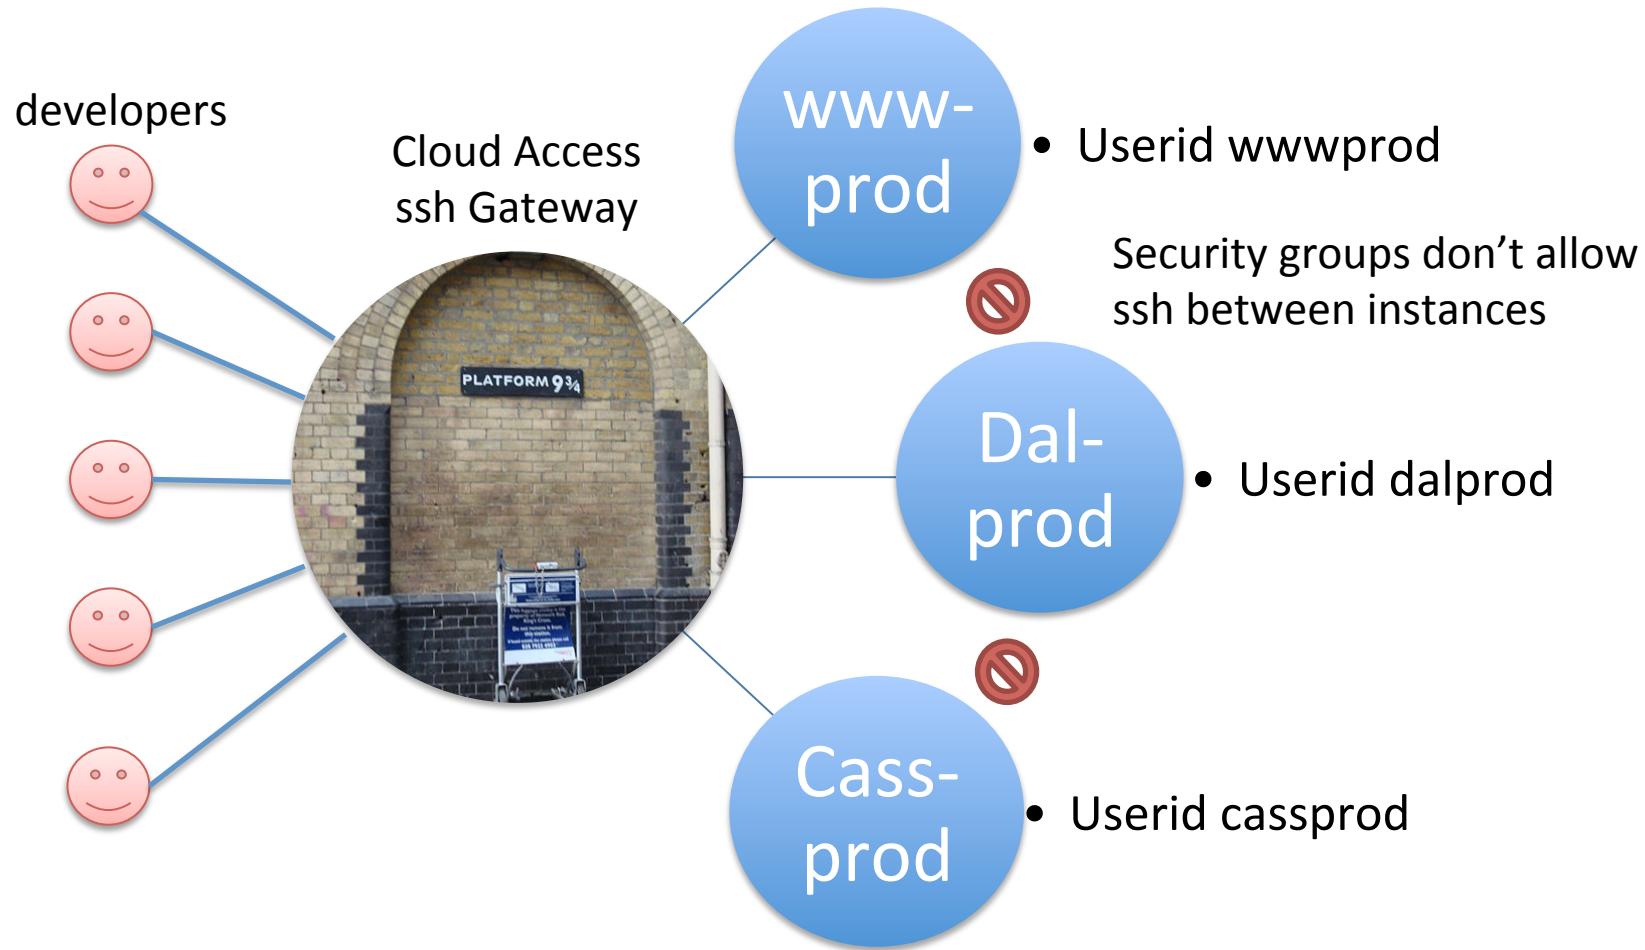
- paastest – for development and testing
 - Fully functional deployment of all services
 - Developer tagged “stacks” for separation
- paasprod – for production
 - Autoscale groups only, isolated instances are terminated
 - Alert routing, backups enabled by default
- paasaudit – for sensitive services
 - To support SOX, PCI, etc.
 - Extra access controls, auditing
- paasarchive – for disaster recovery
 - Long term archive of backups
 - Different region, perhaps different vendor

Reservations and Billing

- Consolidated Billing
 - Combine all accounts into one bill
 - Pooled capacity for bigger volume discounts

<http://docs.amazonwebservices.com/AWSConsolidatedBilling/1.0/AWSConsolidatedBillingGuide.html>
- Reservations
 - Save up to 71% on your baseline load
 - Priority when you request reserved capacity
 - Unused reservations are shared across accounts

Cloud Access Control



Now Add Code

Netflix has open sourced a lot of
what you need, more is on the way...



Netflix Open Source Strategy


- Release PaaS Components git-by-git
 - Source at github.com/netflix – we build from it...
 - Intros and techniques at techblog.netflix.com
 - Blog post or new code every few weeks
- Motivations
 - Give back to Apache licensed OSS community
 - Motivate, retain, hire top engineers
 - “Peer pressure” code cleanup, external contributions


<http://netflix.github.com>

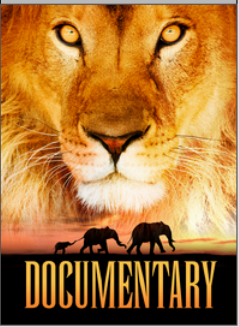
NETFLIX**Netflix Open Source Center**


RepositoriesCommit TimelineMailing Lists


Our Repositories


Astyanax**THRILLERS**


Curator**ACTION & ADVENTURE**


Priam**DOCUMENTARY**


CassJMeter**DRAMA**


Servo**COMEDY**


Aws-Autoscaling**Animal TALES**

Exhibitor**FOREIGN**

Archaius**SCI-FI & FANTASY**

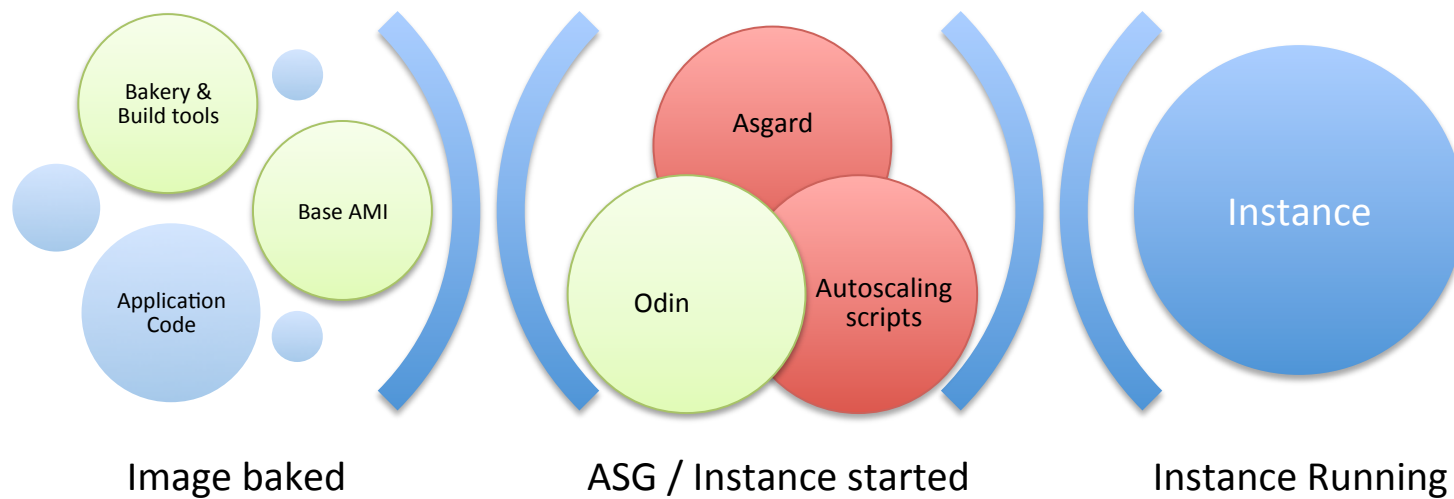
Asgard**MUSIC & MUSICALS**

SimianArmy**CLASSICS**


Eureka**CRIME ACTION**

NETFLIX

Instance creation



Application Launch

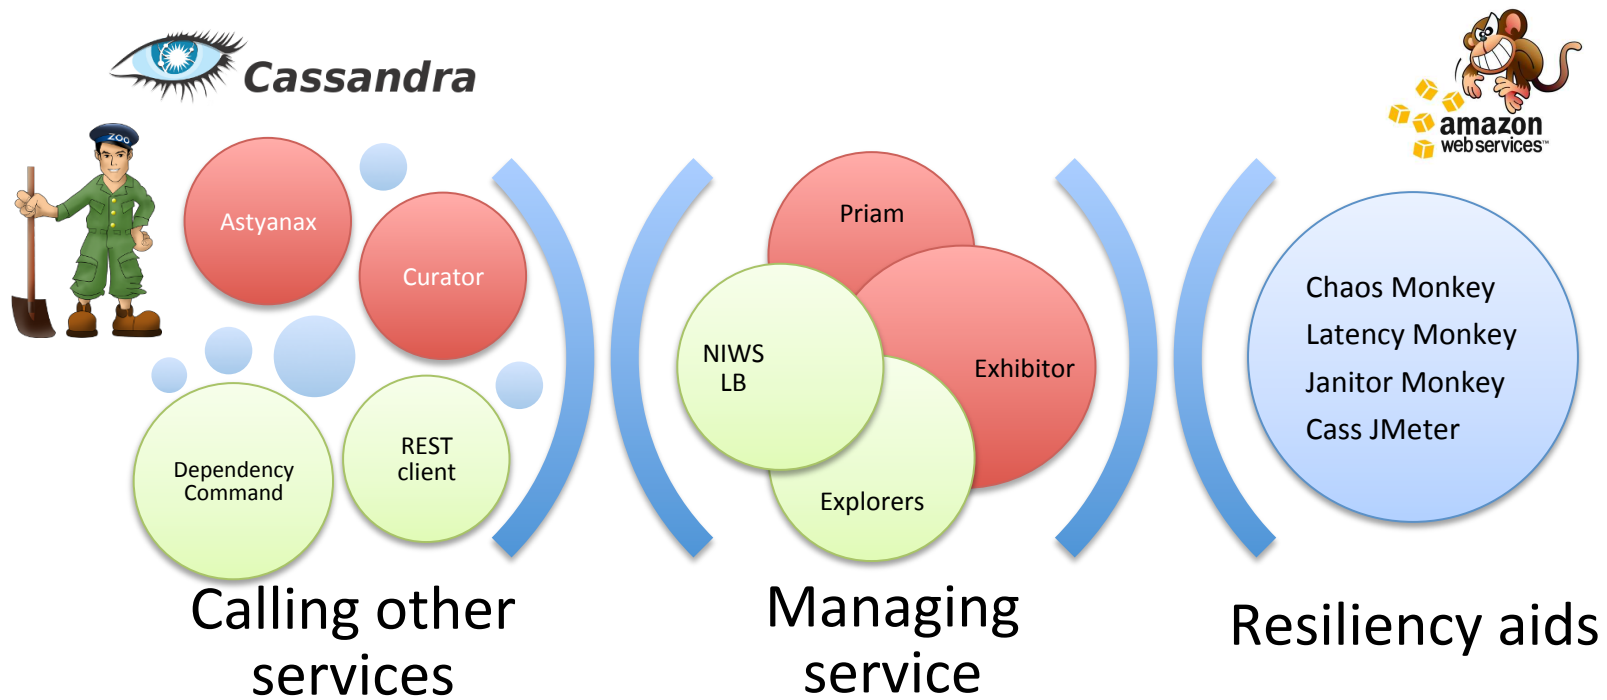


Application initializing

Registering,
configuration



Runtime



Asgard

<http://techblog.netflix.com/2012/06/asgard-web-based-cloud-management-and.html>

The screenshot displays the Asgard web interface for managing a cluster of Sequential Auto Scaling Groups. The browser address bar shows the URL `asgardprod/us-east-1/cluster/show/obiwan`. The interface includes a navigation bar with icons for Home, App, AMI, Cluster, ELB, EC2, SDB, SNS, SQS, RDS, and Task. The main content area is titled "Manage Cluster of Sequential Auto Scaling Groups" and includes a recommendation: "Recommended next step: Switch traffic to the preferred group, then delete legacy group".

Two cluster panels are visible:

- obiwan-v063**: Shows "Launch and Terminate are disabled". It has a "Resize" button set to 9 min / 12 max, and "Delete", "Disable", and "Enable" buttons. Below, it lists "9 instances grouped by state" with a table showing 9 instances in "InService" state, all with a "Build" of 580 and "OUT_OF_SERVICE" status.
- obiwan-v064**: Also shows "Launch and Terminate are disabled". It has a "Resize" button set to 9 min / 12 max, and "Delete", "Disable", and "Enable" buttons. Below, it lists "9 instances grouped by state" with a table showing 9 instances in "InService" state, all with a "Build" of 583 and "UP" status.

On the right, the "Create Next Group" panel for **obiwan-v065** is shown, including fields for AMI Image ID, Instance Type, and Instance Counts (Min: 9, Desired: 9, Max: 12).

Handwritten green annotations highlight key details:

- A green box around the URL bar is annotated with "This cluster contains two ASGs".
- A green box around the "obiwan-v063" table is annotated with "No traffic on old version".
- A green box around the "obiwan-v064" table is annotated with "Live traffic on new version".

NETFLIX

Asgard

- Runs in a VM in our datacenter
 - So it can deploy to an empty account
 - Groovy/Grails/JVM based
 - Supports all AWS regions on a global basis
- Hides the AWS credentials
 - Use AWS IAM to issue restricted keys for Asgard
 - Each Asgard instance manages one account
 - One install each for paastest, paasprod, paasaudit

Discovery and Coordination

- Eureka – Service Directory
 - Map an instance to a service in a private namespace
 - Load balance over clusters of instances
 - Foundation service, first to deploy
- Highly available distributed coordination
 - Netflix Curator for common Zookeeper usage patterns
 - Netflix Exhibitor manages Zookeeper reliably

<http://techblog.netflix.com/2012/09/eureka.html>

<http://techblog.netflix.com/2011/11/introducing-curator-netflix-zookeeper.html>

<http://techblog.netflix.com/2012/04/introducing-exhibitor-supervisor-system.html>

Archaius Dynamic Properties Service

<http://techblog.netflix.com/2012/06/announcing-archaius-dynamic-properties.html>

- Dynamic hierarchical & propagates in seconds
 - Client timeouts, feature set enables
 - Region specific service endpoints
 - etc. etc.
- Used to configure everything
 - So everything depends on it...
 - Pluggable backend storage interface
 - Files, SimpleDB, Cassandra

Open Source Projects

Legend

Github / Techblog

Apache Contributions

Techblog Post

Coming Soon

Priam

Cassandra as a Service

Astyanax

Cassandra client for Java

CassJMeter

Cassandra test suite

Cassandra Multi-region EC2
datastore support

Aegisthus

Hadoop ETL for Cassandra

Explorers

Governator - Library lifecycle
and dependency injection

Odin

Workflow orchestration

Async logging

Exhibitor

Zookeeper as a Service

Curator

Zookeeper Patterns

EVCache

Memcached as a Service

Eureka / Discovery
Service Directory

Archaius

Dynamics Properties Service

EntryPoints

Server-side latency/error
injection

REST Client + mid-tier LB

Configuration REST endpoints

Servo and Autoscaling Scripts

Honu

Log4j streaming to Hadoop

Circuit Breaker

Robust service pattern

Asgard - AutoScaleGroup
based AWS console

Chaos Monkey

Robustness verification

Latency Monkey

Janitor Monkey

Bakeries and AMI

Build dynaslaves

NETFLIX

Roadmap for 2012

- More resiliency and improved availability
- More automation, orchestration
- “Hardening” the platform, code clean-up
- Lower latency for web services and devices
- IPv6 – now running in prod, rollout in process
- More open sourced components
- See you at AWS Re:Invent in November...

Takeaway

Netflix has built and deployed a scalable global Platform as a Service.

Key components of the Netflix PaaS are being released as Open Source projects so you can build your own custom PaaS.

<http://github.com/Netflix>
<http://techblog.netflix.com>
<http://slideshare.net/Netflix>

<http://www.linkedin.com/in/adriancockcroft>
<http://www.linkedin.com/in/ruslanmeshenberg>

@adrianco @rusmeshenberg #netflixcloud



Amazon Cloud Terminology Reference

See <http://aws.amazon.com/> This is not a full list of Amazon Web Service features

- AWS – Amazon Web Services (common name for Amazon cloud)
- AMI – Amazon Machine Image (archived boot disk, Linux, Windows etc. plus application code)
- EC2 – Elastic Compute Cloud
 - Range of virtual machine types m1, m2, c1, cc, cg. Varying memory, CPU and disk configurations.
 - Instance – a running computer system. Ephemeral, when it is de-allocated nothing is kept.
 - Reserved Instances – pre-paid to reduce cost for long term usage
 - Availability Zone – datacenter with own power and cooling hosting cloud instances
 - Region – group of Avail Zones – US-East, US-West, EU-Eire, Asia-Singapore, Asia-Japan, SA-Brazil, US-Gov
- ASG – Auto Scaling Group (instances booting from the same AMI)
- S3 – Simple Storage Service (http access)
- EBS – Elastic Block Storage (network disk filesystem can be mounted on an instance)
- RDS – Relational Database Service (managed MySQL master and slaves)
- DynamoDB/SDB – Simple Data Base (hosted http based NoSQL datastore, DynamoDB replaces SDB)
- SQS – Simple Queue Service (http based message queue)
- SNS – Simple Notification Service (http and email based topics and messages)
- EMR – Elastic Map Reduce (automatically managed Hadoop cluster)
- ELB – Elastic Load Balancer
- EIP – Elastic IP (stable IP address mapping assigned to instance or ELB)
- VPC – Virtual Private Cloud (single tenant, more flexible network and security constructs)
- DirectConnect – secure pipe from AWS VPC to external datacenter
- IAM – Identity and Access Management (fine grain role based security keys)

