Leaping into "The Cloud"

Rewards, Risks, and Mitigations

Ken Johnston, Principal Group Program Manager, Bing
Seth Eliot, Senior Knowledge Engineer, Test Excellence

Microsoft

Better Software West – June 13, 2012

About Us

Seth

- Microsoft Engineering Excellence: Best practices for services and cloud
- Bing: Massive, distributed, data processing service
- Microsoft ExP: Data Driven Decision Making
- Amazon.com: Video, Music, and Kindle eBook services

Ken

- Principal Group Program Manager, Bing
- Office 2010, MSN, Hosted Exchange
- Director of Test Excellence

What Do You Know?

- Just beginning with cloud?
- Who has a major project coming up?
- Who has already implemented a cloud service?
- Anything ever gone wrong?

Introduction

- About Clouds
- Cloud Rewards
- Getting Into The Cloud
- 5 Amazing Cloud Case Studies
  - Rewards, Risks & Mitigations
- Testing in The Cloud

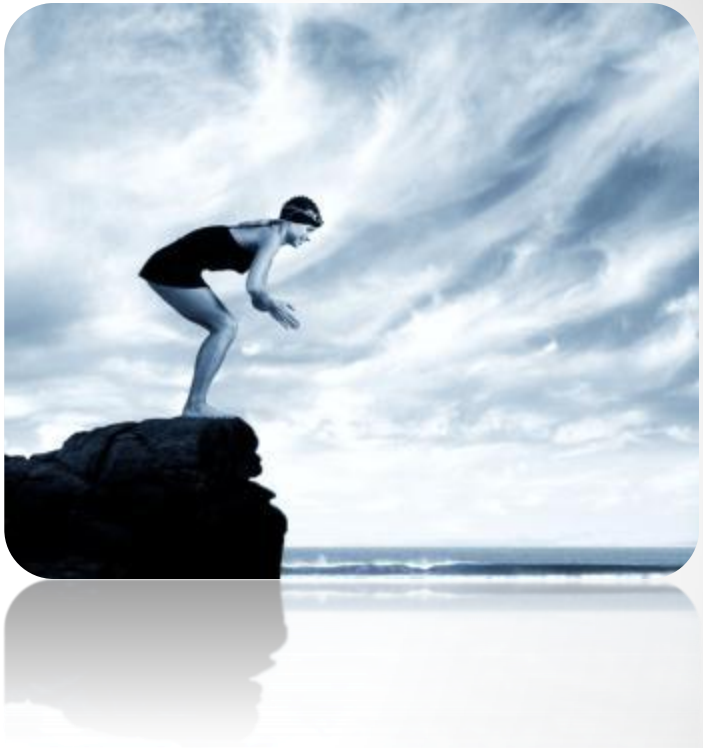

The latest version of this slide deck can be found at:
http://www.setheliot.com/blog/bsc-west-2012/

About Clouds

Three Ingredients of The Cloud

1. Standardized IT capability or service
   - No customizing for each customer
   - Economies of Scale - rote, repeatability

2. Pay Per Use
   - The power of zero

3. Self-Service Deployment
   - Fully Automated
   - DevOps, NoOps

[Staten, 2010]

The Cloud's Secret Sauce

Virtualization

Elasticity

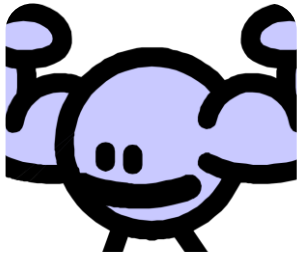

Power

Happiness

Automatically.....?

Yes, Like This…

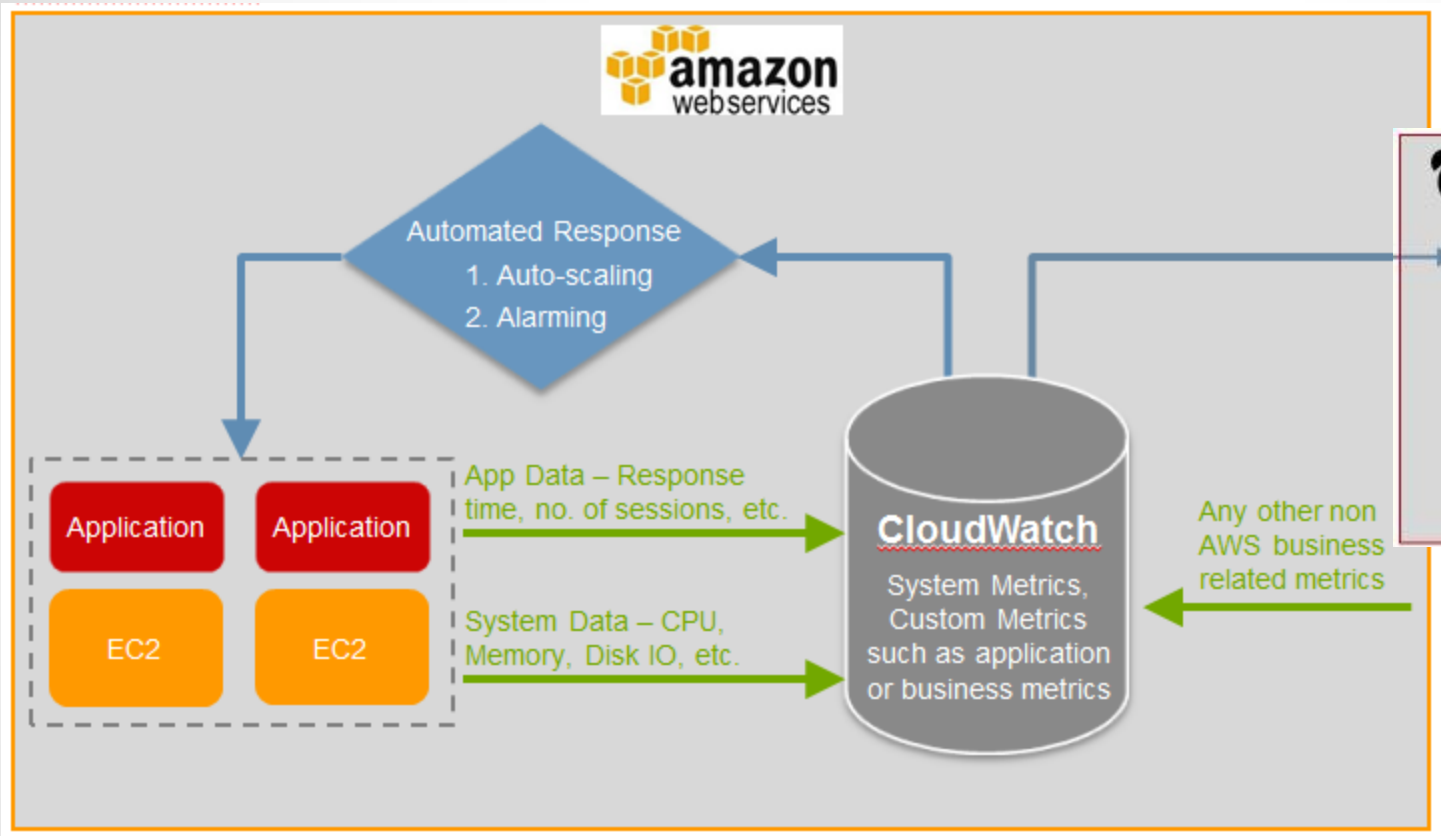

[Netflix Autoscaling, 2012]

- Scale Up: alarm at 75% of target threshold with a 5-10 minute delay before automated action takes place
- Scale Down: slowly, using time as a proxy to avoid removing capacity too quickly

Three Layers of of Clouds

Cloud Category	The Cloud handles...	Examples
Cirrus	16,500 to 40,000 ft	
Altocumulus	6,500 to 23,000 ft.	
Cumulus	Surface to 10,000 ft	

Three Layers of of Clouds

Cloud Category	The Cloud handles…	Examples
SaaS: Software	e.g., Office Application Functionality	Microsoft Office Web Apps Google Docs
PaaS: Platform	Relational Database Management Systems	Microsoft SQL Azure Amazon RDS
	Frameworks and Runtimes	Microsoft Windows Azure - .NET Google App Engine – Java, Python
	Messaging Queue	Microsoft Azure Queue Amazon SQS
IaaS: Infrastructure	Servers	Amazon EC2 - Linux, Windows Rackspace Cloud Servers - Linux
	Storage	Amazon S3 / SDB - BLOB / Table Microsoft Windows Azure Storage
	CDN	Windows Azure CDN Amazon CloudFront
	Network	Amazon Virtual Private Cloud

XaaS = …as a Service

Are Clouds for Real?

Unparalleled Market Growth

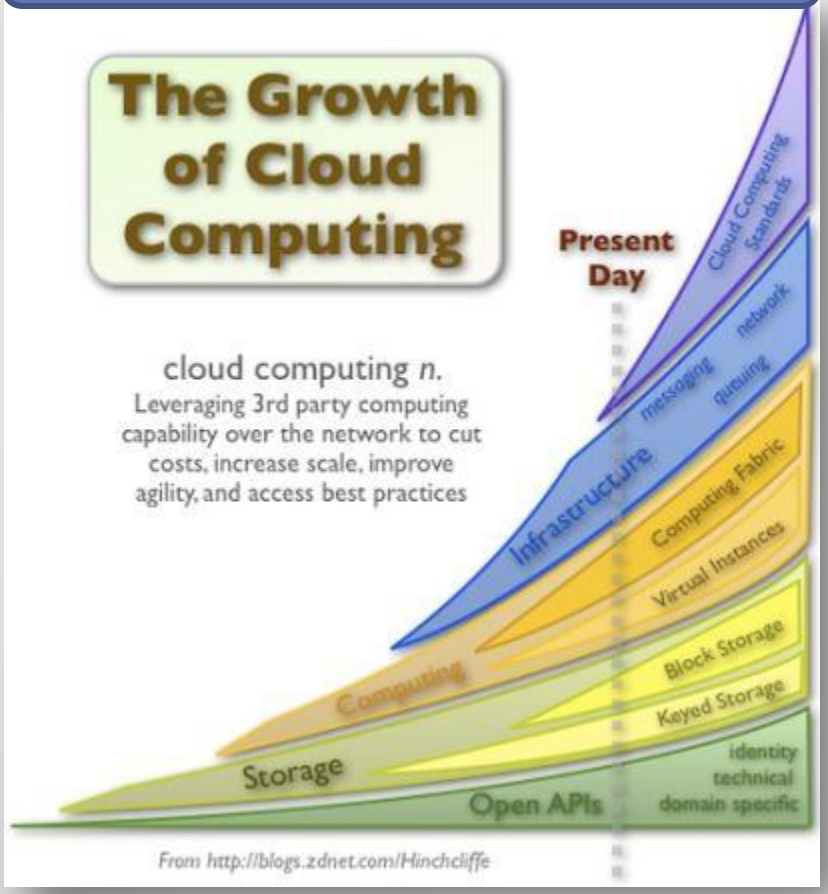

[Hinchcliffe, 2009]

Massive Adoption

- Global cloud computing to grow from $37.8 billion 2010 to **$121.1 billion in 2015**

[R&M, 2010]

- By **2015**, business revenues from IT innovation enabled by the cloud could reach **US$1.1 trillion** a year

[Microsoft, March 2012]

14 Million New Jobs by 2015

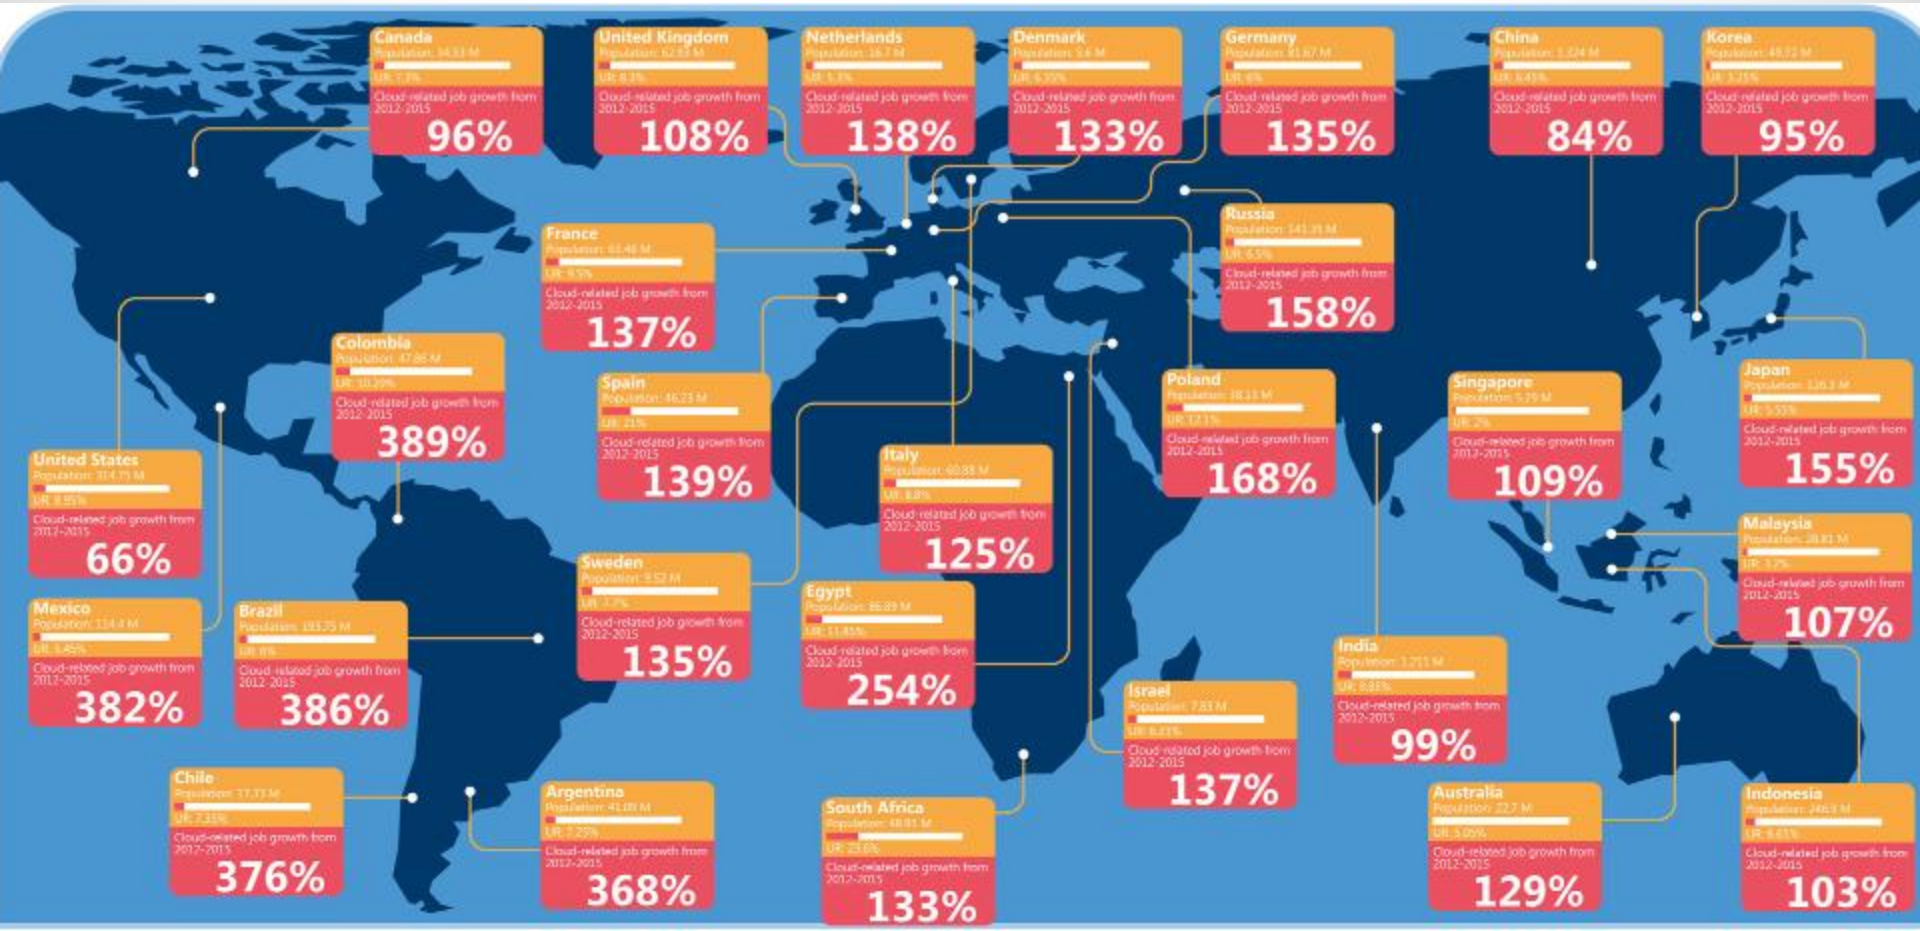

[Microsoft, March 2012]

Really? Are Clouds for Real?

- Massive Investments
  - Cloud To Command 90% of Microsoft's R&D Budget [Forbes, 2011]
    - **~8.6 Billion in 2011**
- Amazing Growth
  [Amazon Growth, 2011]
- Steep competition
  - 90 Cloud Computing Companies to Watch in 2011
    [CCJ, 2011]

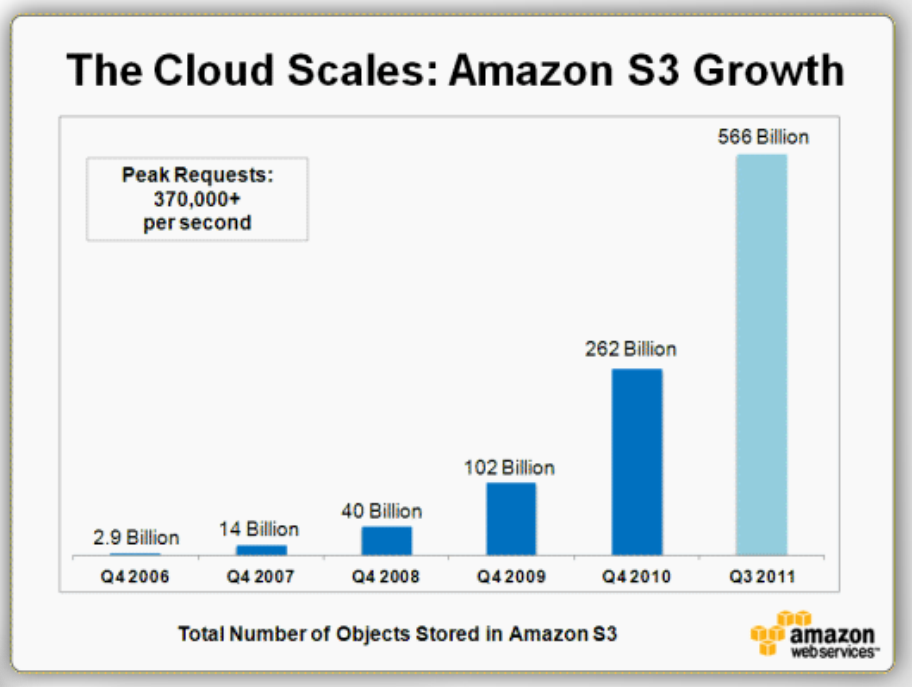

amazon web services

salesforce

vmware

Zuora

rackspace HOSTING

HOSTING.com

Appistry

rPath

Joyent

Cloud Rewards

The Promise of the Cloud

Promises, Promises…

The Cloud Makes Many Promises

You are Empowered to Leverage These

- You Have an active role

Cloud Promise + Your Actions = Rewards

Rewards, A 40,000 ft. View

1. On demand capacity	Elasticity
2. Lower Cost	The Cloud is your data center
3. Disaster Recovery	Backups
4. Fault tolerance	Redundancy
5. Ease of management	Automation and APIs
6. Rewards Guaranteed	SLA – Service Level Agreement
7. Easy Integration	Many Services - One Provider

2. Lower Cost - The Cloud is your data center

- Asset Utilization
  - Data center server utilization averages 5%-20% [Berkeley 2009]
- Hardware Costs
  - Data center performance - only increases with additional investment.
- Power Efficiency
  - Power Usage Effectiveness (PUE) for Data Center
  - Industry average 2.0
  - Microsoft Chicago:1.22
  - Microsoft Quincy 1.15 [Microsoft DC, 2011]

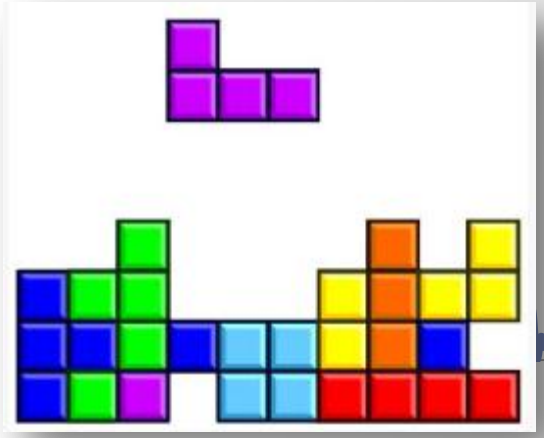

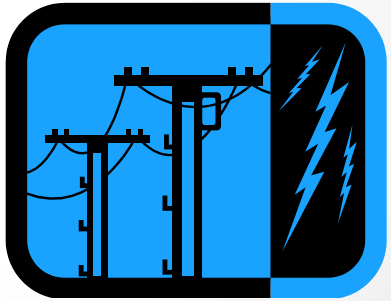

Continued….

Lower Cost - The Cloud is your data center (cont)

- Security
  - Network security devices
  - Security software licenses
  - Staffing
  - Regulatory compliance
  - Physical security requirements
- Supply Chain Management
  - Ordering servers and components costs money and time
- Personnel
  - Operating data centers
  - Scaling and managing physical growth

Economies of Scale

Larger datacenters have almost 50% lower TCO per server

MAIN DATA CENTER COST BUCKETS

- **Server hardware costs** (~45% of total costs)
- **Facility & operations** (~25%)
- **Hardware labor costs** (~15%)
- **Power costs** (~15%)

ANNUAL TCO/SERVER DECLINES W/SCALE

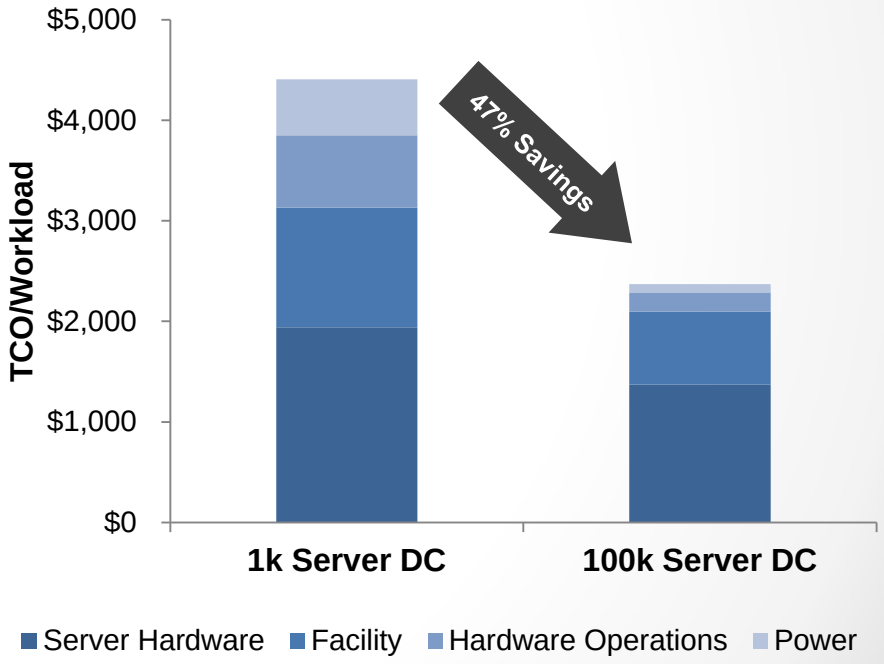

Source: Microsoft

3. Disaster Recovery & 4. Fault Tolerance

Service Robustness Enabled by The Cloud

- Multiple, smaller servers for **Redundancy**
- Handle load spikes via **Elastic Scalability**
- **Backups** leverage IaaS storage
- Use the tools via API – **Automate**

But how about when clouds turn stormy?

5. Ease of management - Automation and APIs

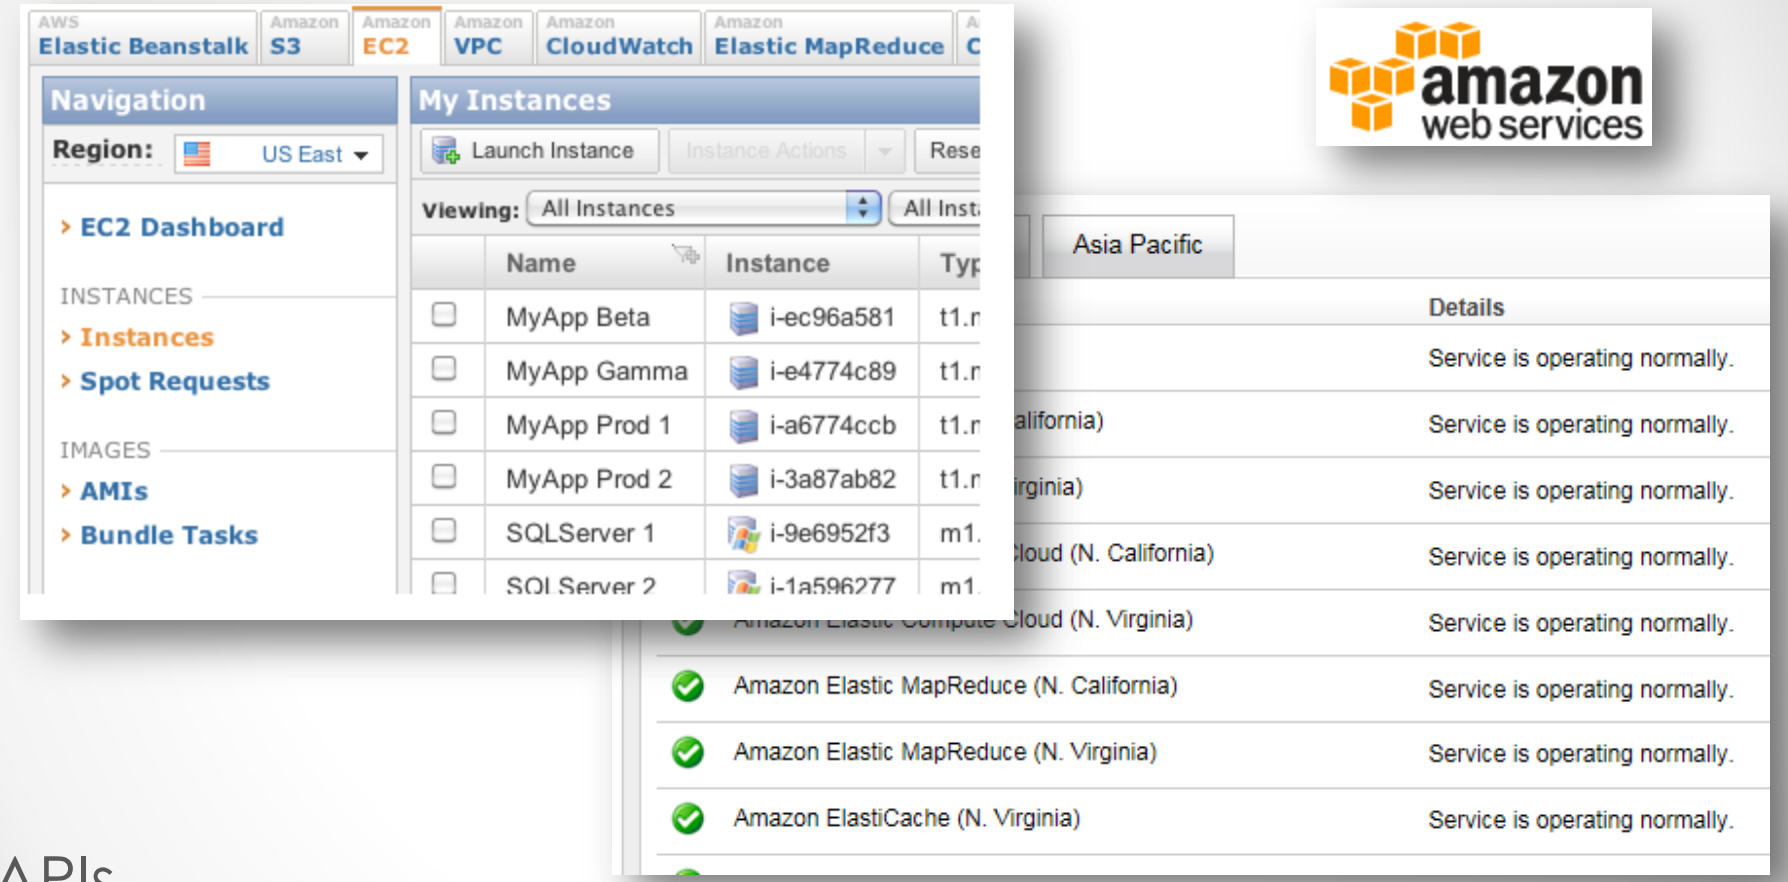

APIs

- Configure Instances, Load Balancers.. Everything
- Monitor via Amazon CloudWatch

New Azure Portal

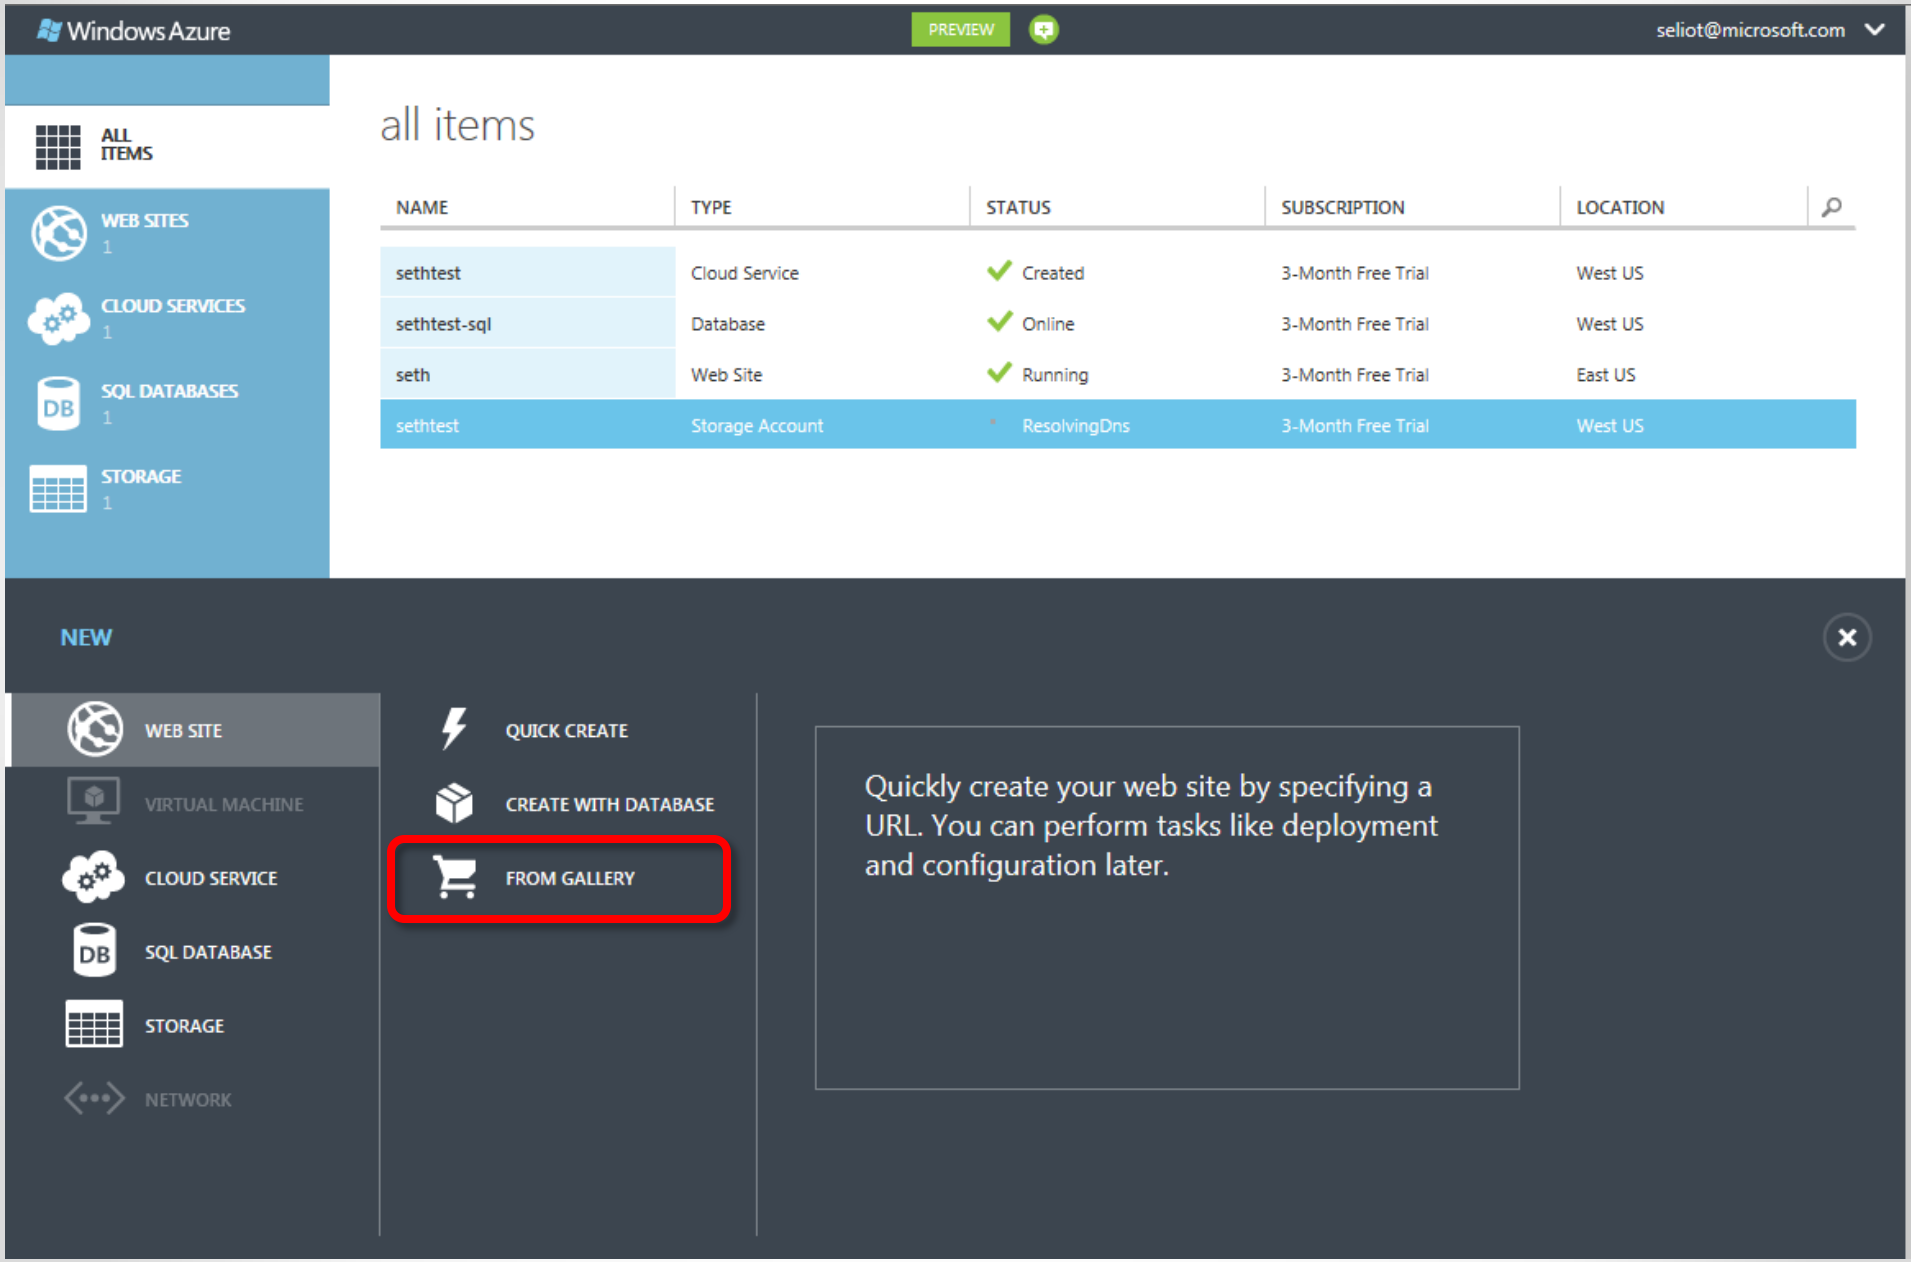

ADD WEB APP

Find Apps for Azure

ALL
BLOGS
CMS
ECOMMERCE
FORUMS

A-Z


Drupal Commerce Kickstart


Joomla! 2.5


mojoPortal


Orchard CMS


Umbraco CMS 5


WordPress

WordPress

WordPress is a state-of-the-art publishing platform with a focus on aesthetics, web standards, and usability.

VERSION	3.3.1
SIZE	4264
RELEASE DATE	1/2/2012
PUBLISHER	WordPress

2

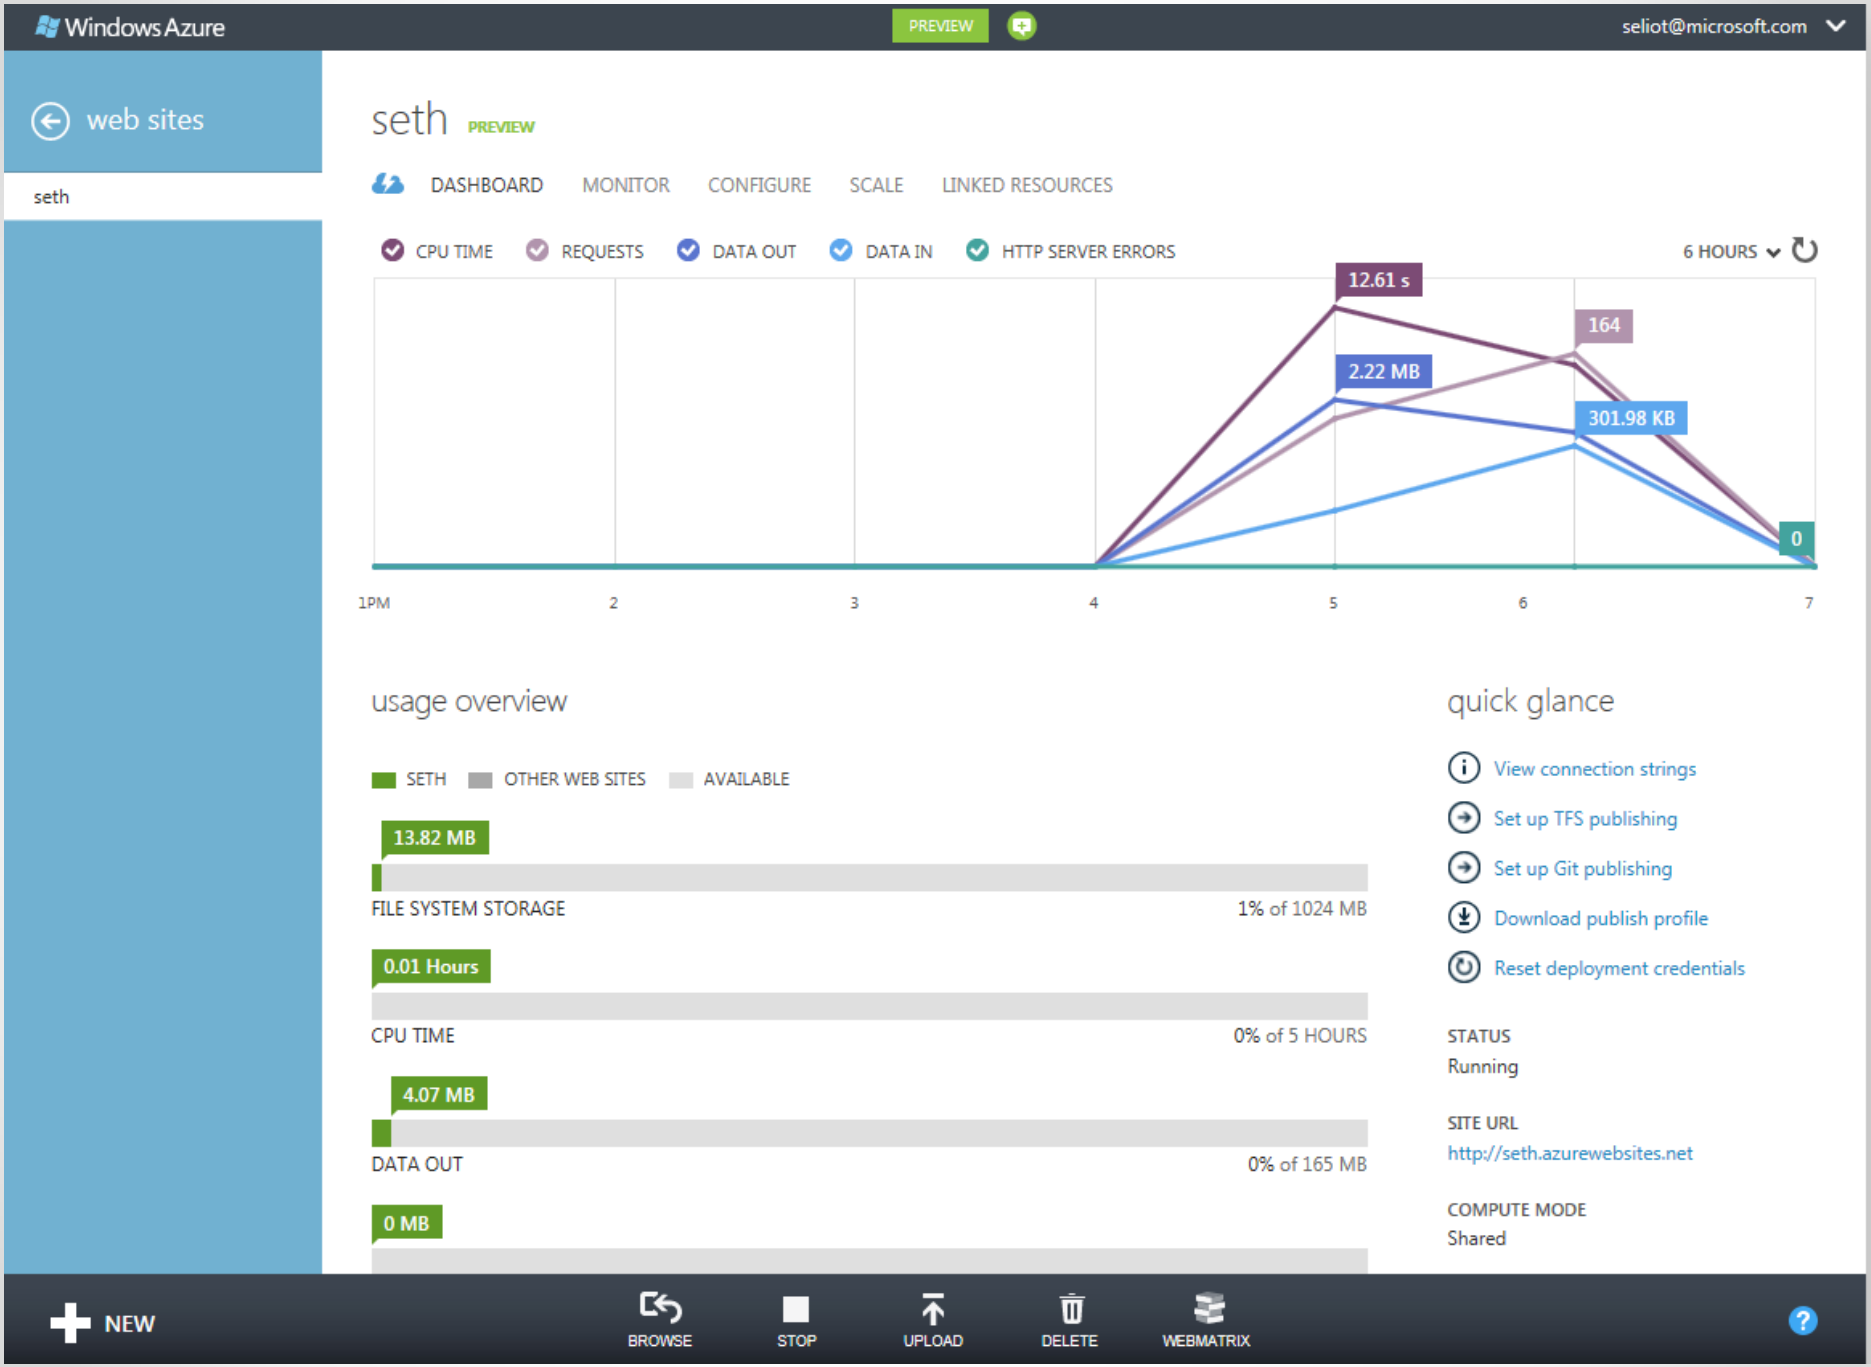
Windows Azure
PREVIEW
seliot@microsoft.com
web sites
seth
seth PREVIEW
DASHBOARD
MONITOR
CONFIGURE
SCALE
LINKED RESOURCES
CPU TIME
REQUESTS
DATA OUT
DATA IN
HTTP SERVER ERRORS
6 HOURS
12.61 s
164
2.22 MB
301.98 KB
0
1PM
2
3
4
5
6
7
usage overview
SETH
OTHER WEB SITES
AVAILABLE
13.82 MB
FILE SYSTEM STORAGE
1% of 1024 MB
0.01 Hours
CPU TIME
0% of 5 HOURS
4.07 MB
DATA OUT
0% of 165 MB
0 MB
quick glance
View connection strings
Set up TFS publishing
Set up Git publishing
Download publish profile
Reset deployment credentials
STATUS
Running
SITE URL
http://seth.azurewebsites.net
COMPUTE MODE
Shared
NEW
BROWSE
STOP
UPLOAD
DELETE
WEBMATRIX

Manage Azure via REST APIs

Windows Azure

Operations on Hosted Services

- List Hosted Services
- Create Hosted Service
- Update Hosted Service
- Delete Hosted Service
- Get Hosted Service Properties
- Create Deployment
- Get Deployment
- Swap Deployment
- Delete Deployment
- Change Deployment Configuration
- Update Deployment Status
- Upgrade Deployment
- Walk Upgrade Domain
- Reboot Role Instance
- Reimage Role Instance
- Rollback Update Or Upgrade
- Check Hosted Service Name Availability
- Get Package

Operations on Sto
Operations on
Operations on
Operations on
Operations for
Operations for
Operations for
Operations on
Operations for
Operations on
Operations on
Operations on
Operations on

6. Rewards Guaranteed - Cloud SLAs

	Microsoft	Amazon	Rackspace	Google
Service	Azure Compute	EC2	Cloud Servers	Apps for Business
SLA	99.9% 99.95%[1]	99.95%	100%	99.9%
Service Credit	10%-25%	10%	5%-100%	3-15 days
Storage	Azure Storage	S3	Cloud Files	
SLA	99.9%	99.9%	99.9%	
Service Credit	10%-25%	10%-25%	10%-100%	

1. If two or more role instances in different fault and upgrade domains [Cloud SLAs]

Example: Azure Storage Uptime = $100\% - \frac{Failed\ Storage\ Transactions}{Total\ Storage\ Transactions}$

- Failures Transactions includes completed but too slow

SLAs, What are They Good For?

- Service Credits will likely not compensate for lost business and negative customer impact.
- Providers pay out service credits, but the cost in publicity is more.
  - The market will reward those that **keep their SLAs**
  - But Enterprise cloud users cannot afford to bet on the wrong provider.
- 99.9% uptime = 9 hrs/yr down
- Must architect defensively
  - More when we get to case studies

7. Easy Integration - Many Services, One Provider

Your Application, plus:

- Storage
- Databases
- Web Servers
- CDN

….all in the Cloud

A Video Download Service

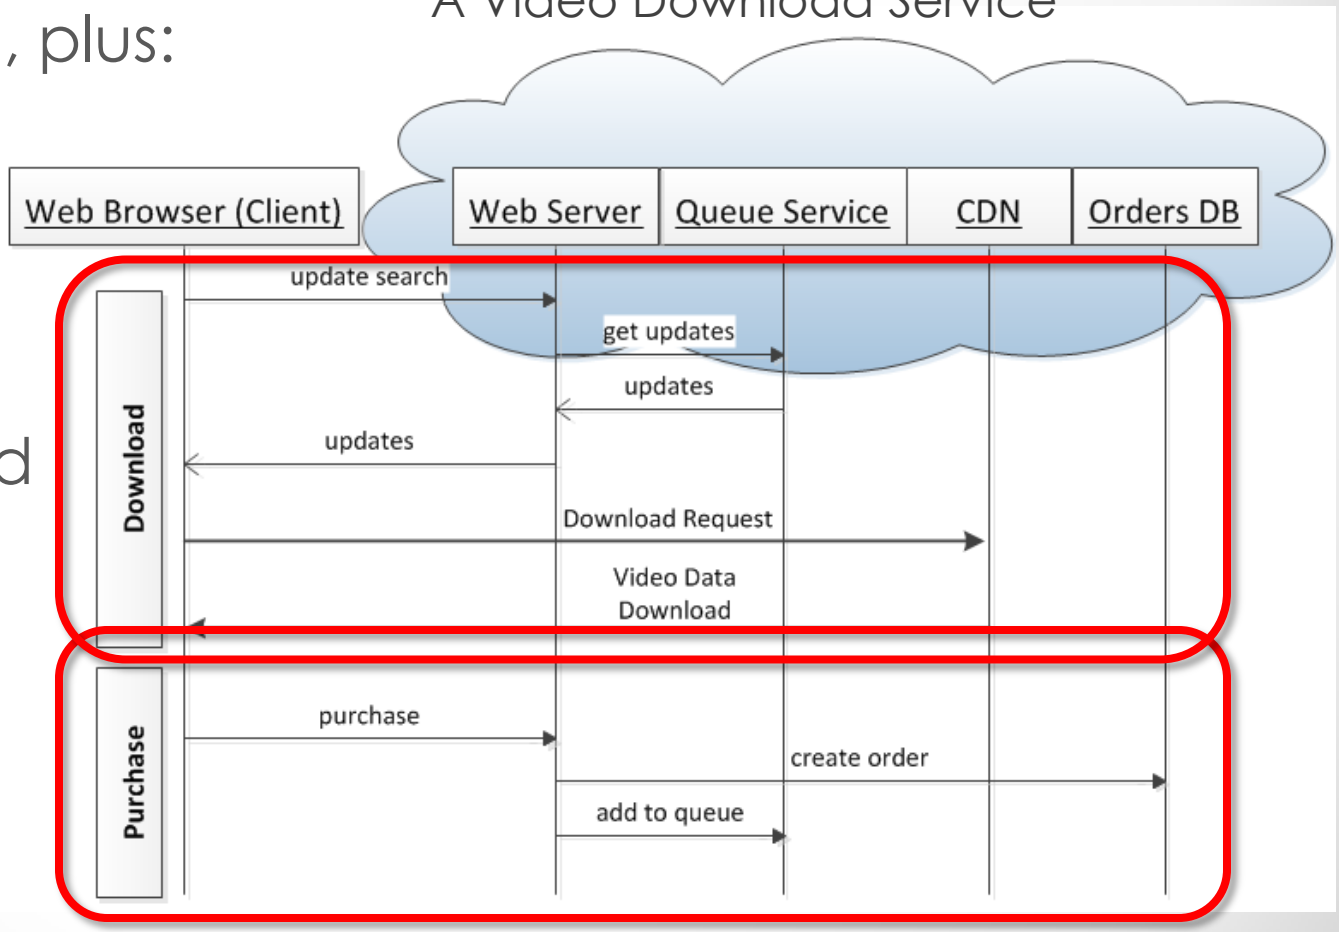

Availability and Interoperability within a single cloud provider

- Simpler than building full solution.

Getting Into The Cloud

Plan Pick and Execute

Plan Your Cloud Migration

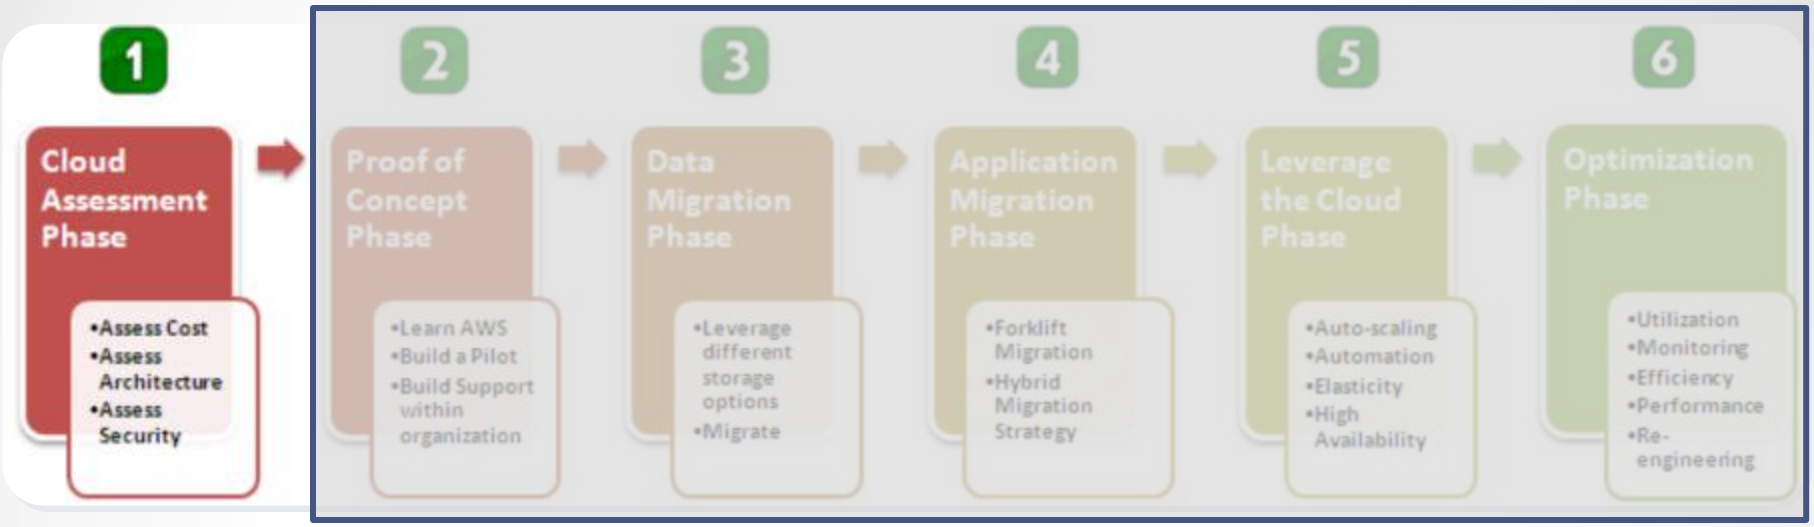

- Model courtesy of Amazon
  - Six step model
  - Plan, proof of concept, execution, optimize
- Leaping into the Cloud is mostly about planning and execution

[AWS Whitepaper]

Plan for each Application

The cloud providers ***want*** you there

Microsoft Azure

- Microsoft Assessment and Planning (MAP) Toolkit [MAP Toolkit]
  - Automatically finds your web apps, web servers and DBs

Estimates what you need

Azure compute instances

SQL Azure DBs

Bandwidth

Storage

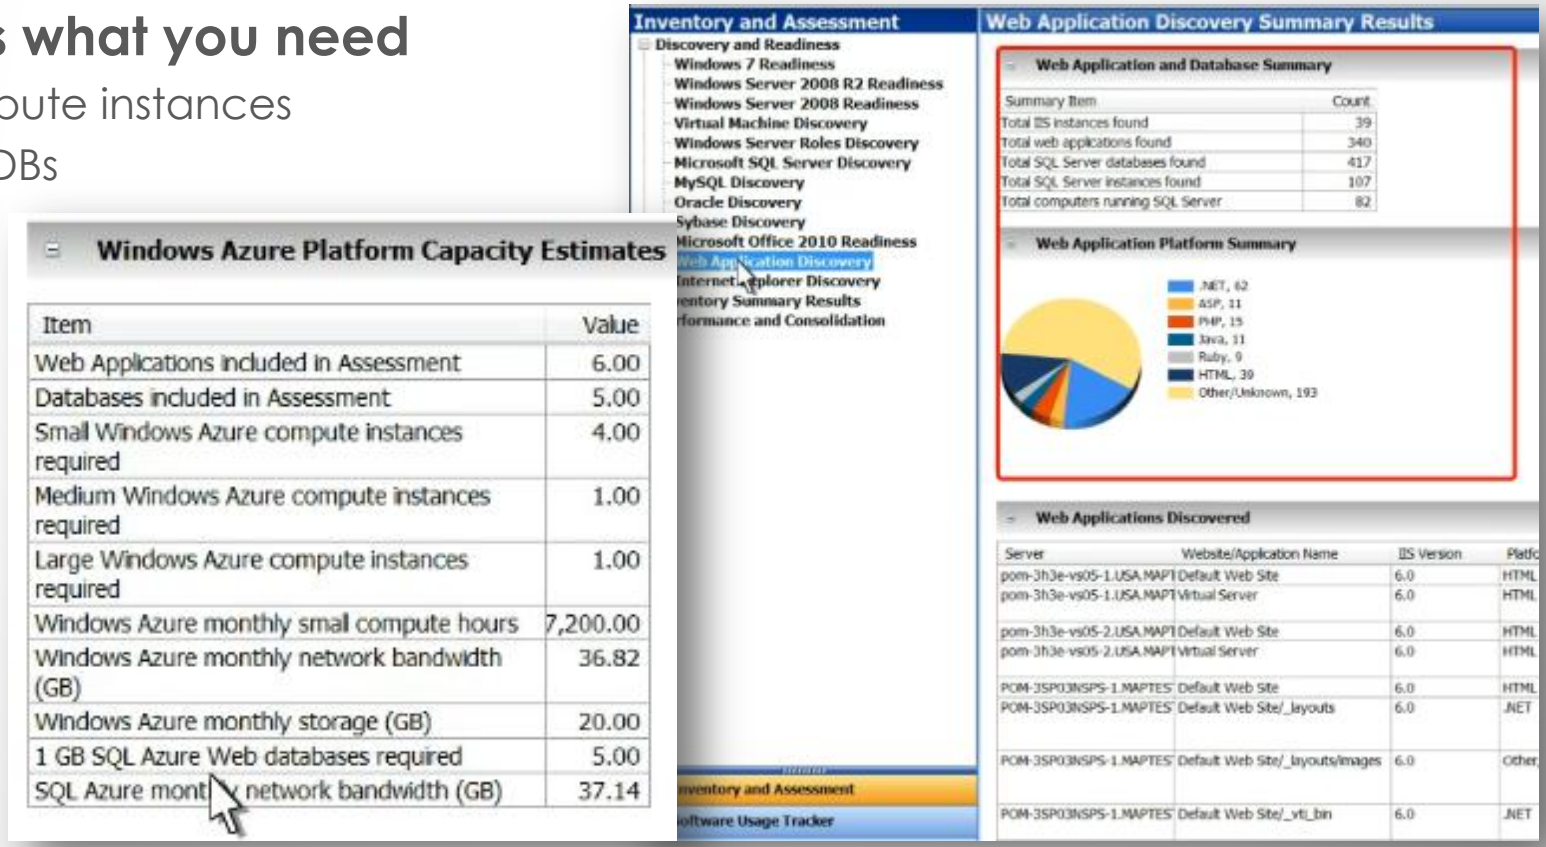

Execute you plan

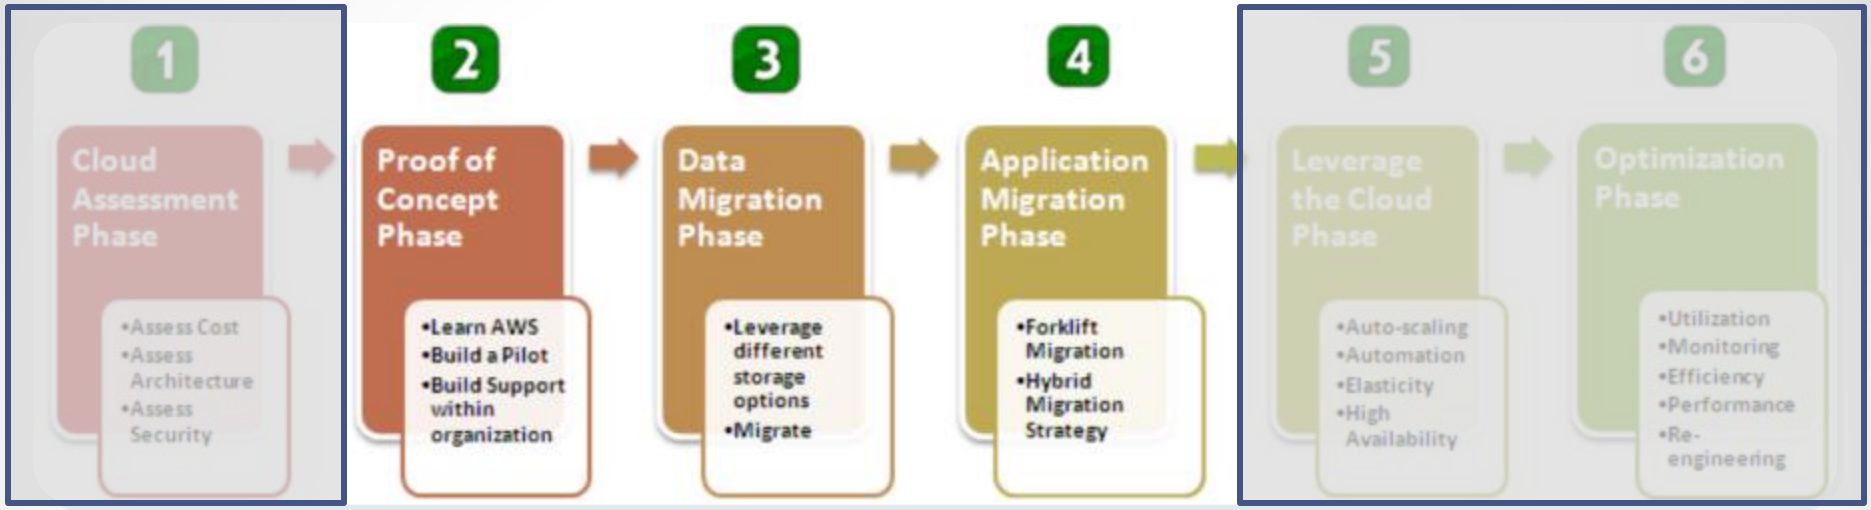

- Proof of Concept
  - Build a trial version in the cloud
  - Plan for data Migration and App Migration
- To do this, you will need to pick a cloud provider

[AWS Whitepaper]

Pick the Services you need

- Types of Services you need (Window/Linux)
- Type of Contract
  - Different pricing
  - Different SLAs
- Security Levels
  - FISMA Compliant – Federal Information Security Management Act [FISMA, 2002]
  - Other Security compliance

Dynamic Pricing Models

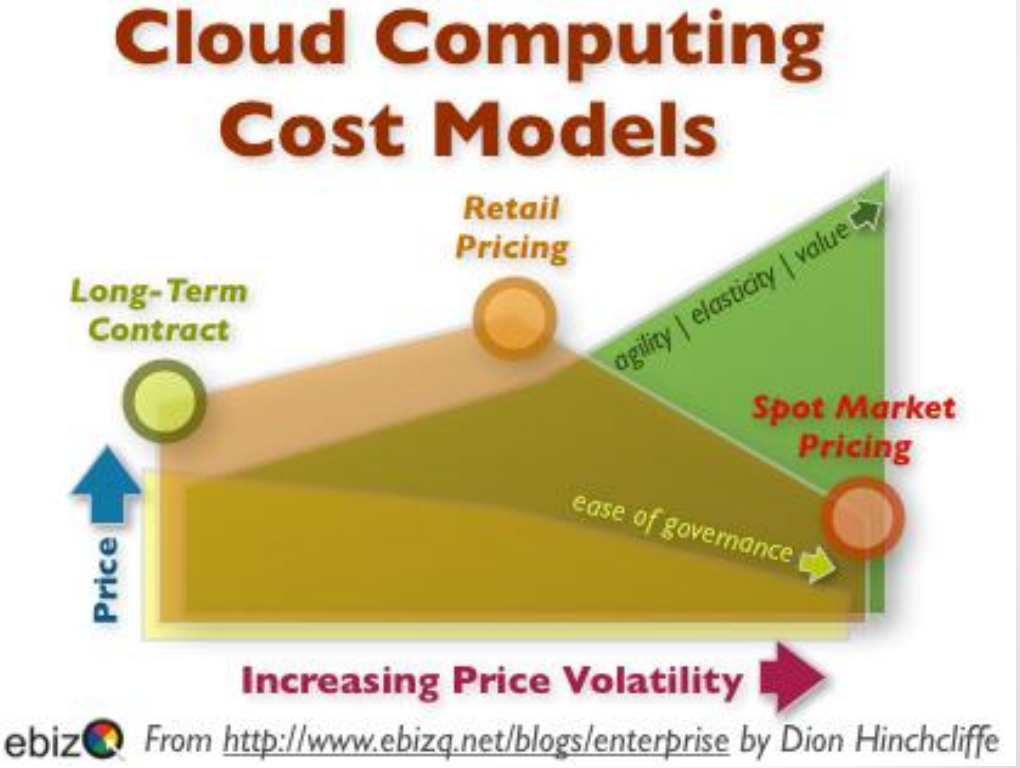

Pick the Right Cloud Provider

- Handy Cloud Computing Price Comparison Engines
  [Cloud Tweaks, 2011]

1. FindTheBest.com
2. ServDex.com
3. CloudSurfing.com
4. Cloudarade.com

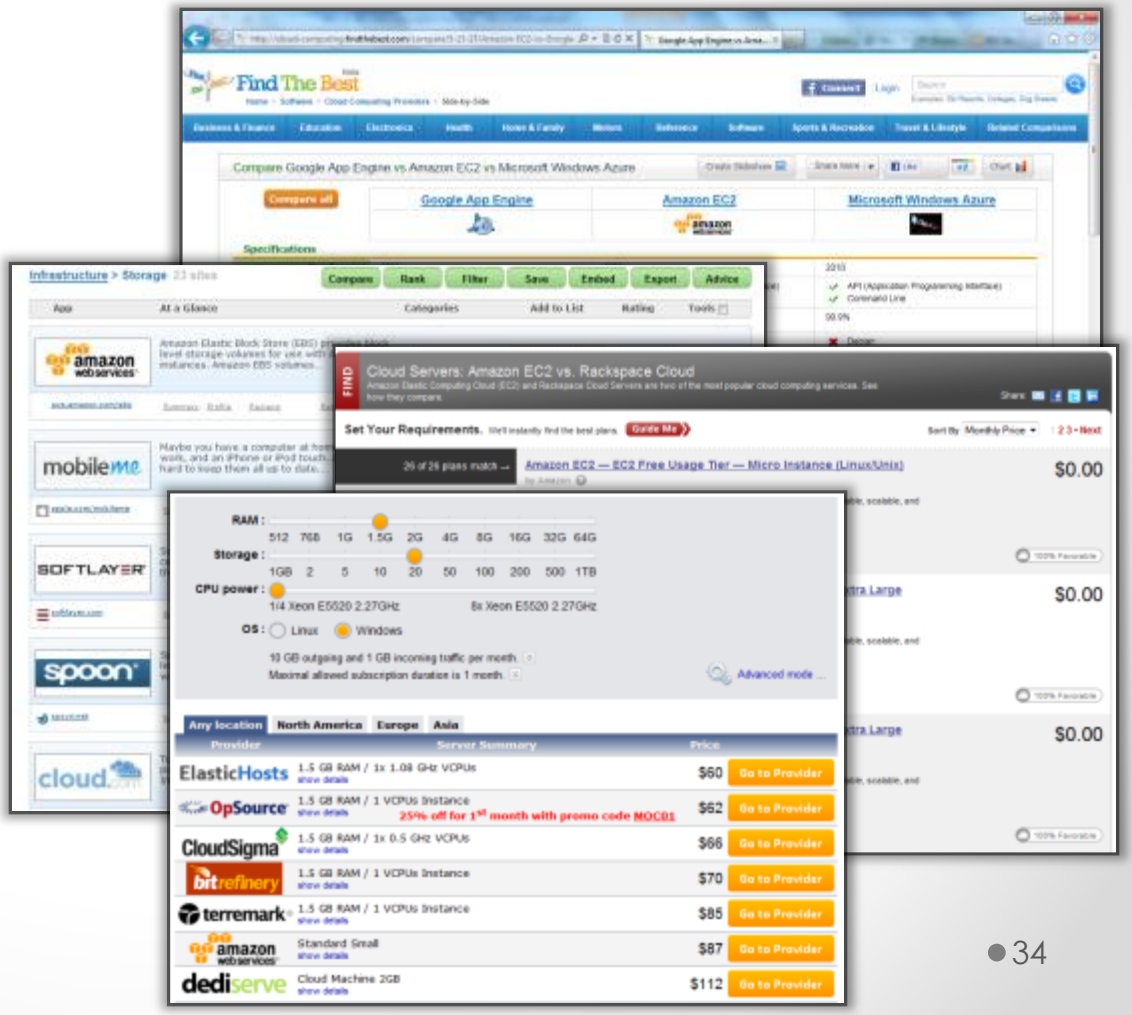

5 Amazing Cloud Case Studies

Rewards, Risks & Mitigations

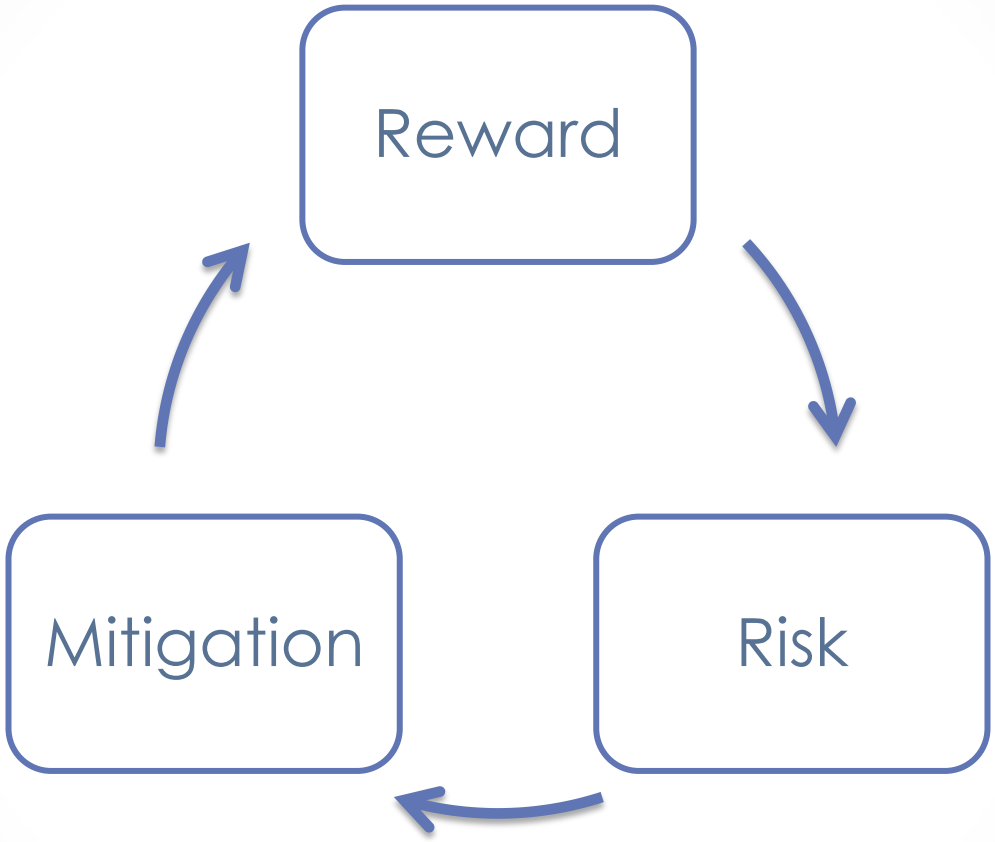
Reward
Risk
Mitigation

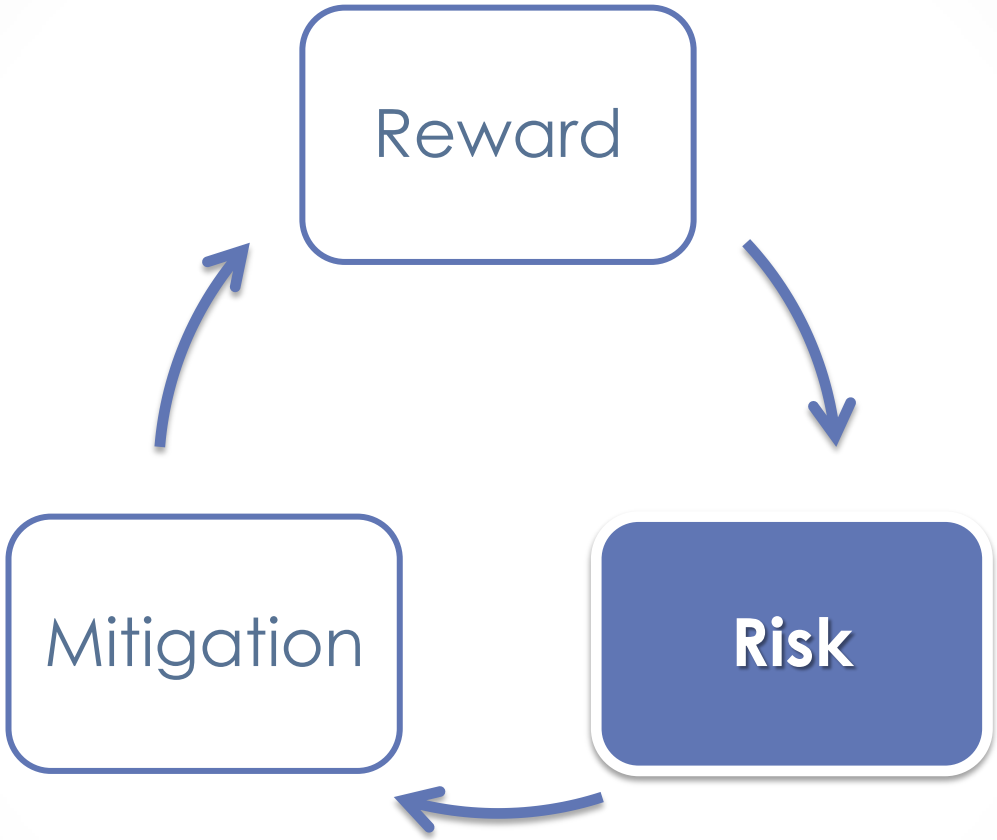
Reward
Risk
Mitigation

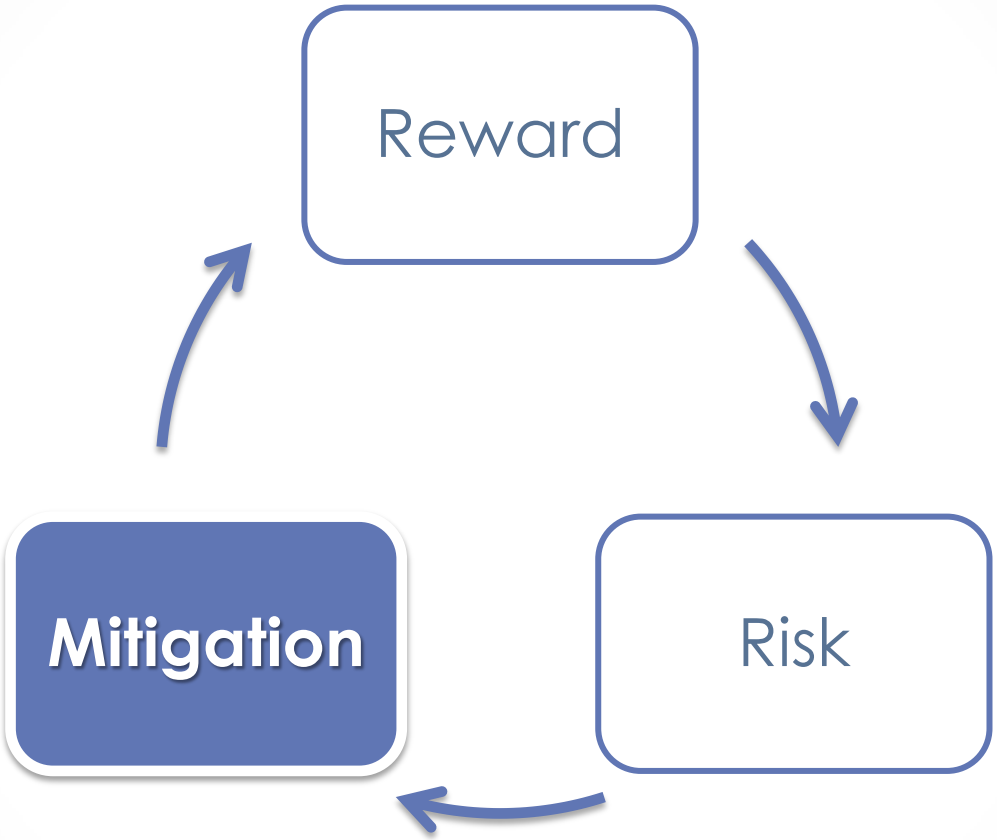
Reward
Risk
Mitigation

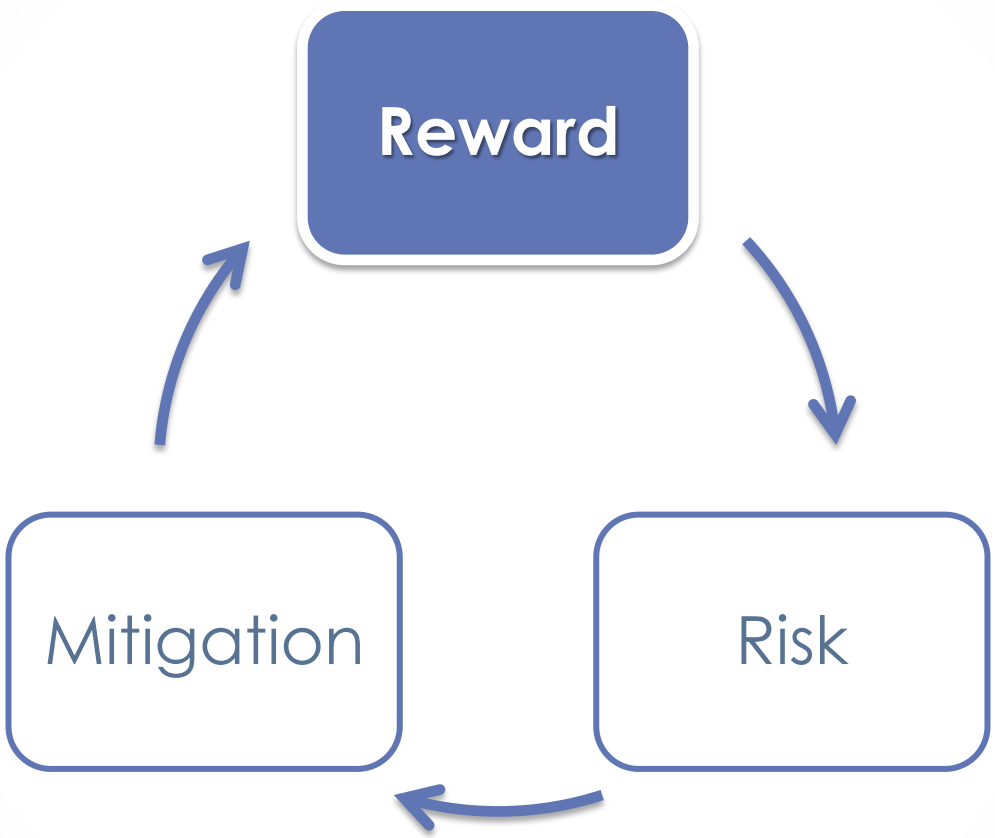
Reward
Risk
Mitigation

Amazon.com Elasticity and Cost Savings

[Jenkins, 2011]

Website Traffic is Spikey

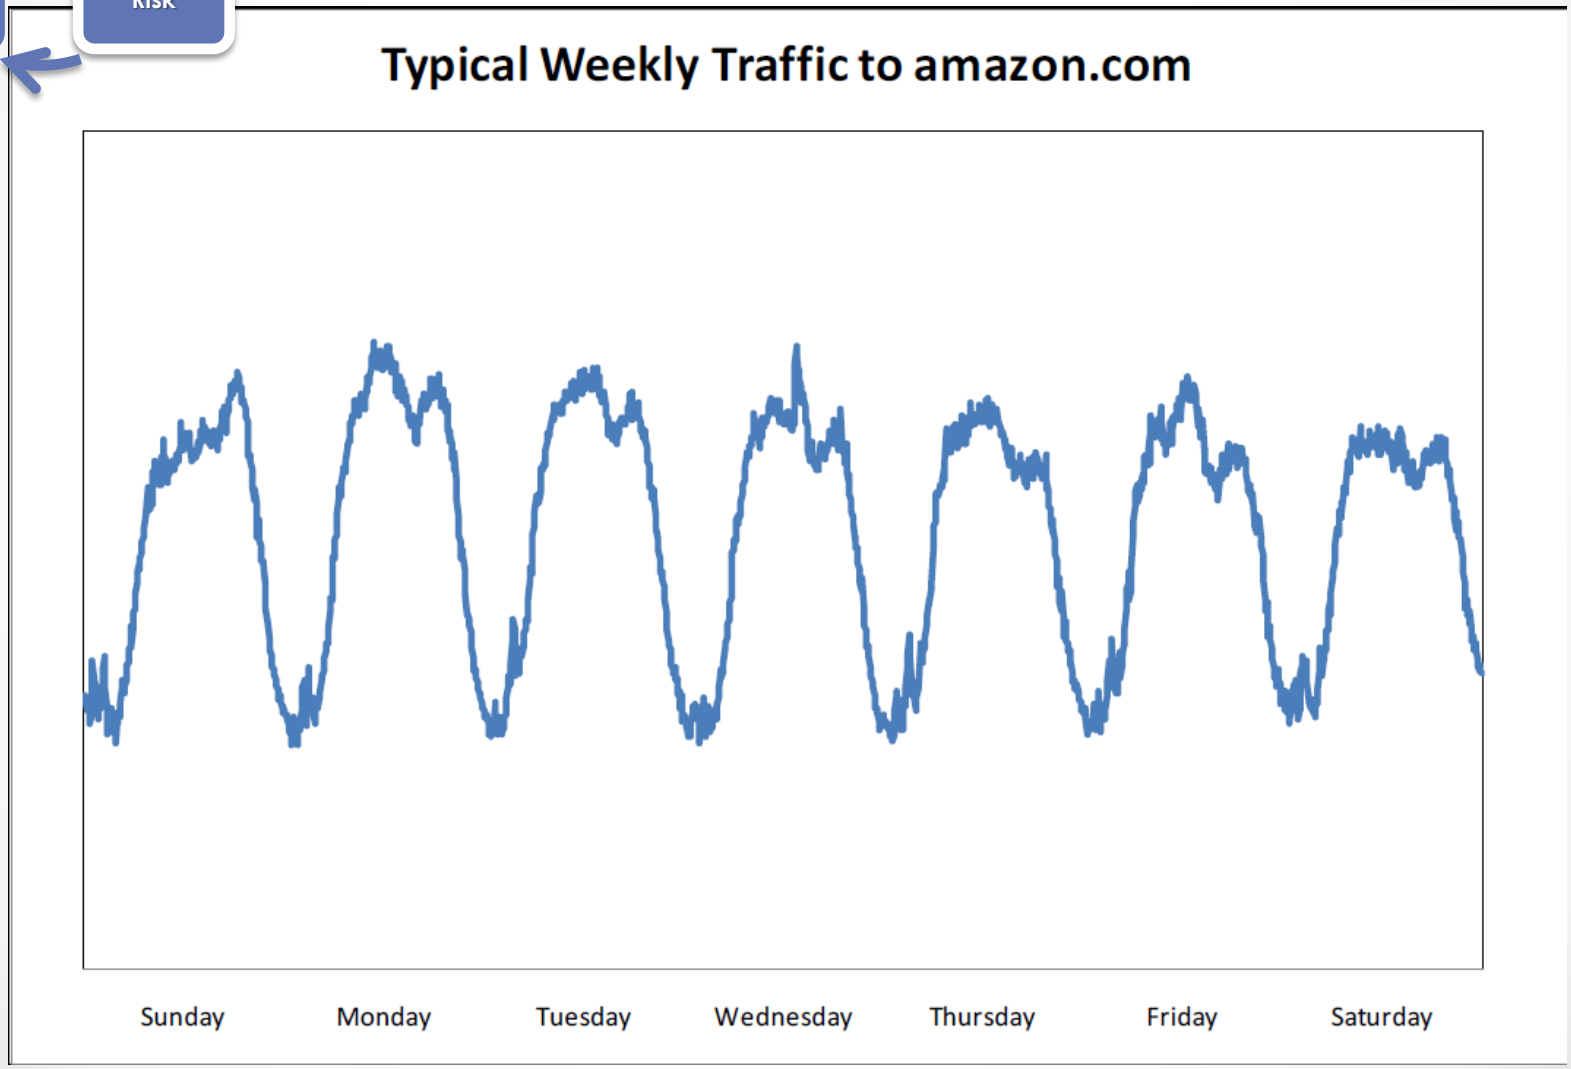


Reward
Risk
Mitigation

Add a Buffer

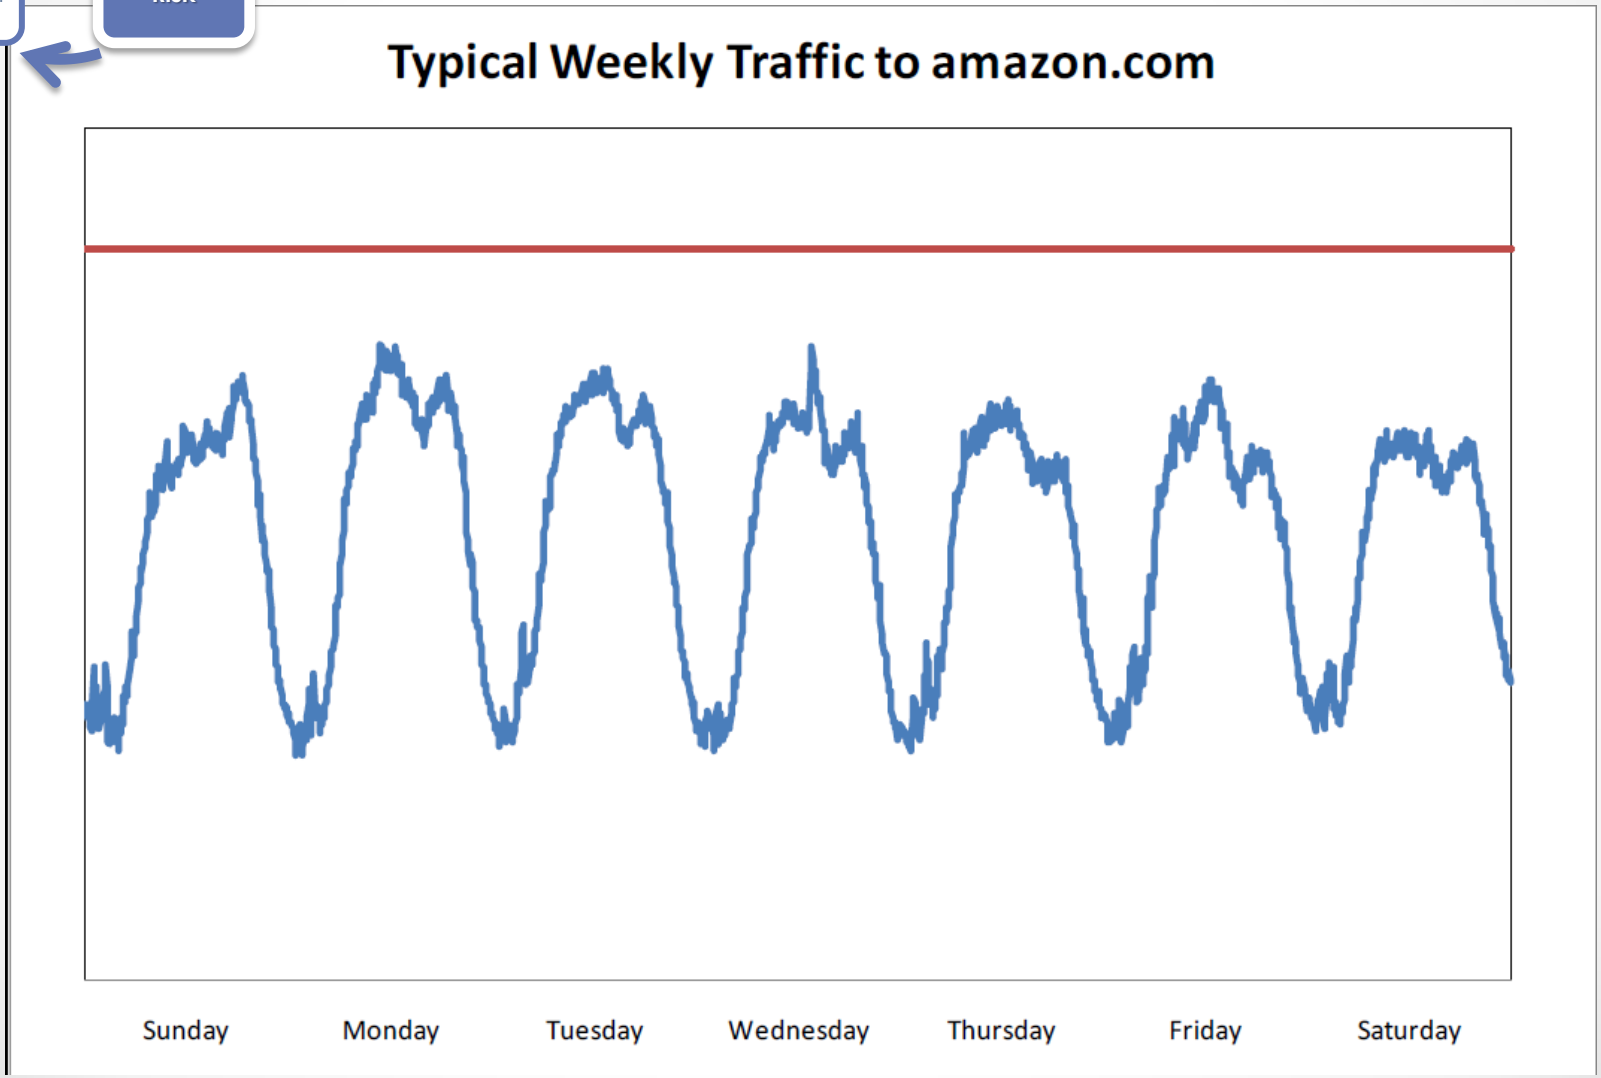
Typical Weekly Traffic to amazon.com
Sunday
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday


Reward
Risk
Mitigation

Major Waste!

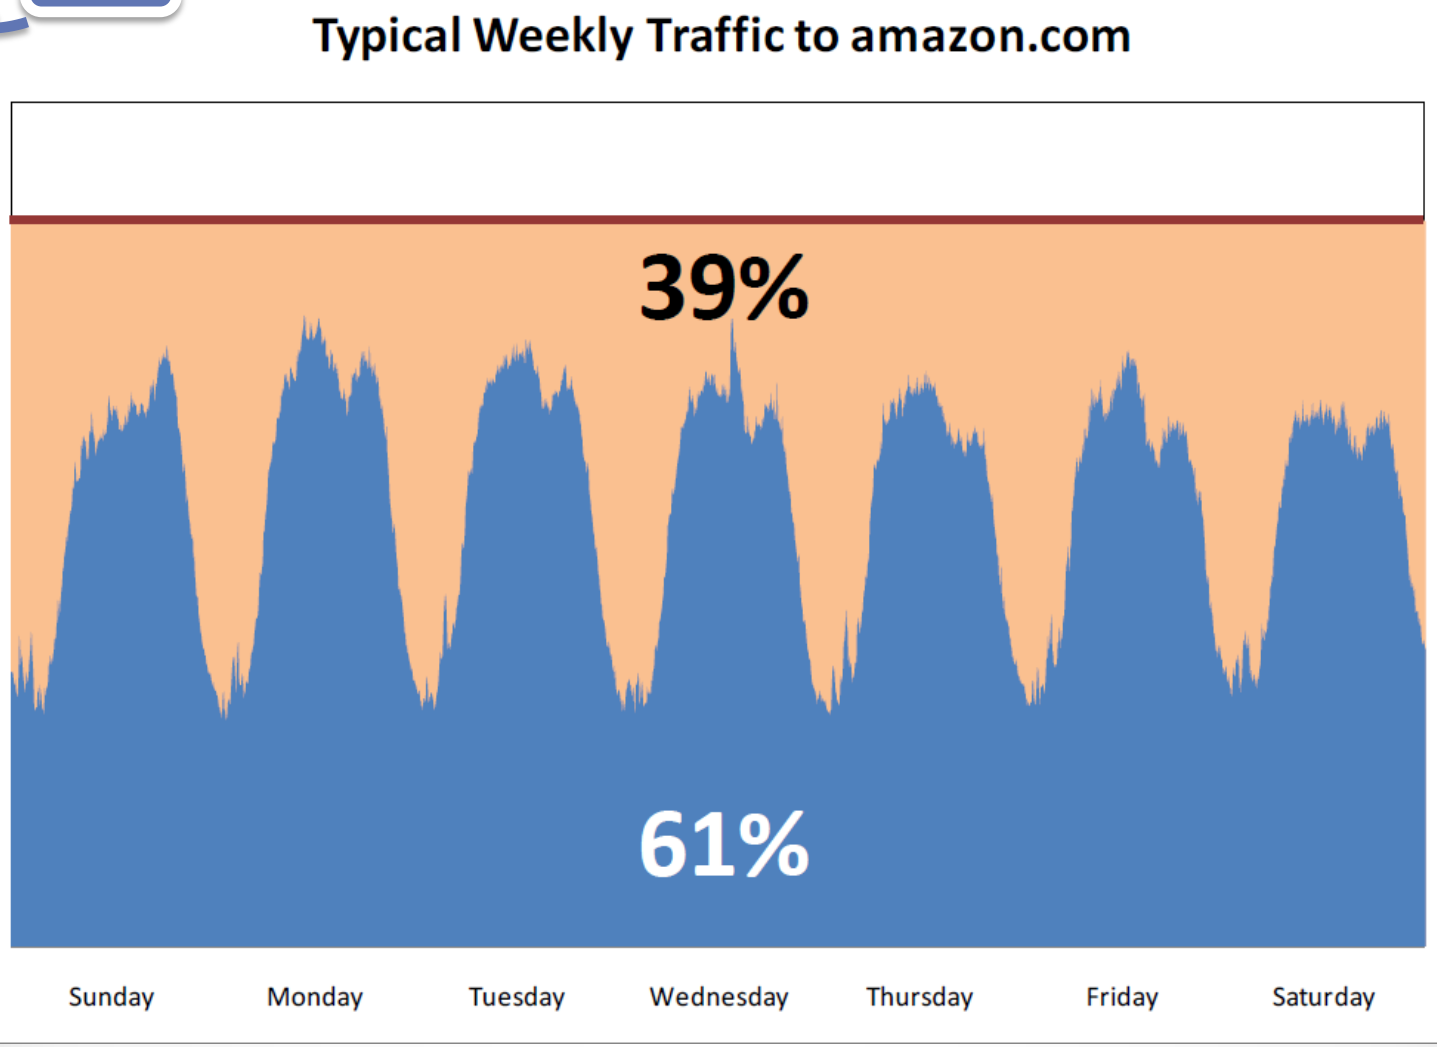
Typical Weekly Traffic to amazon.com
39%
61%
Sunday
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday

But it's Even Worse

Seasonality Spikes

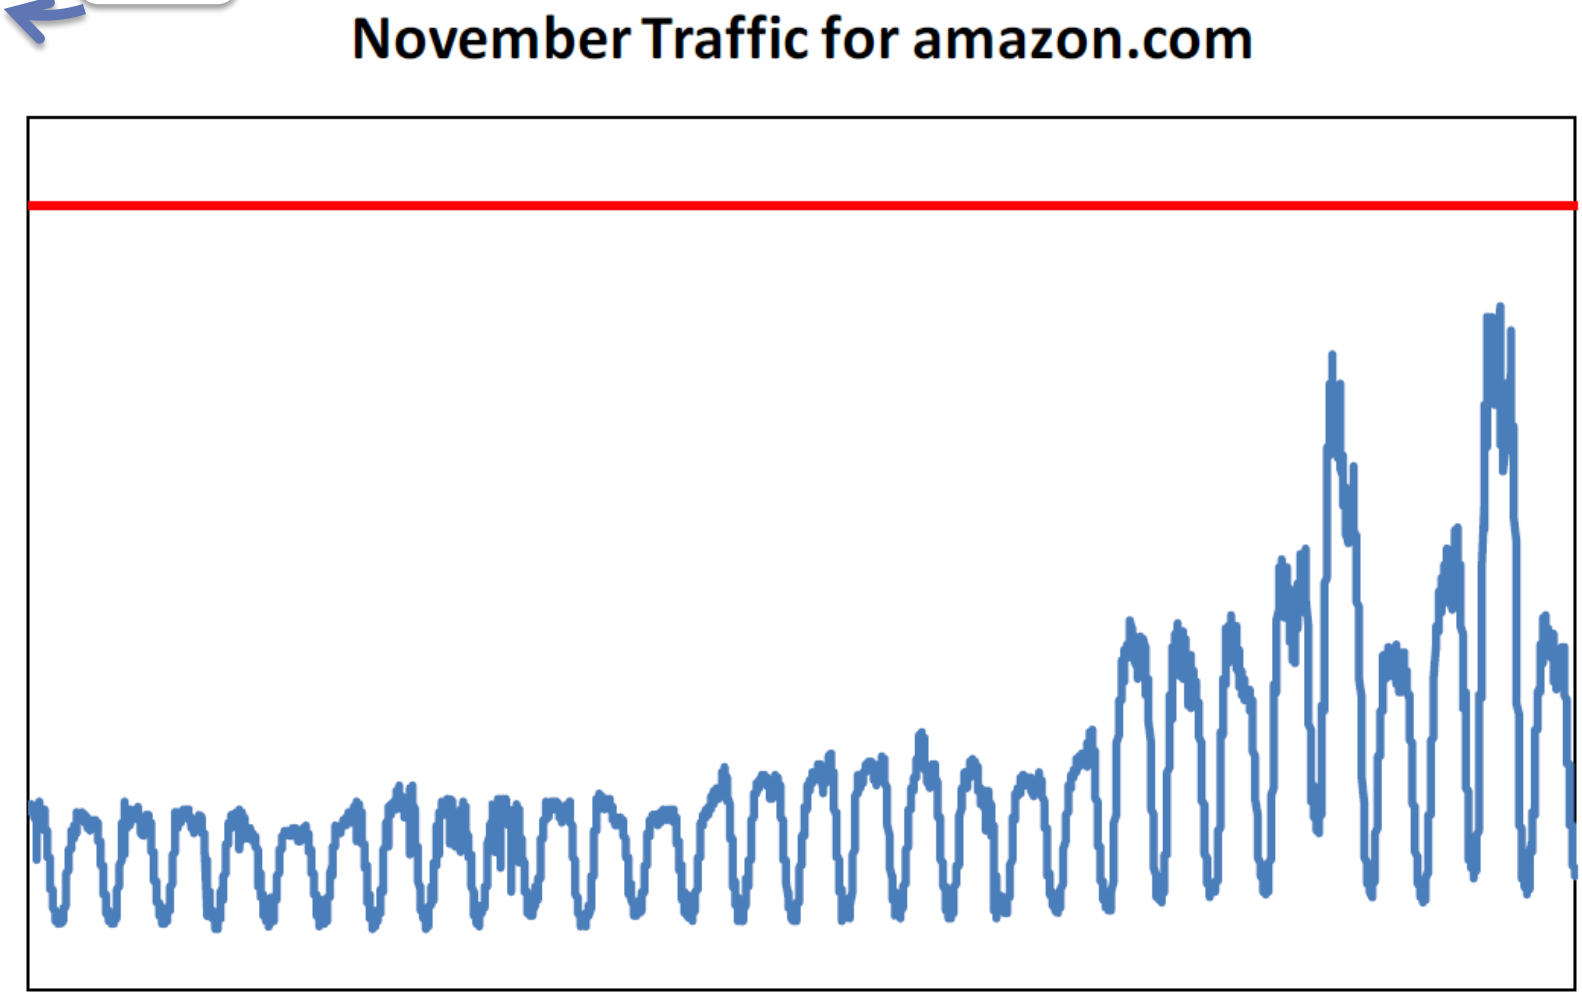


Reward
Risk
Mitigation

Big Waste

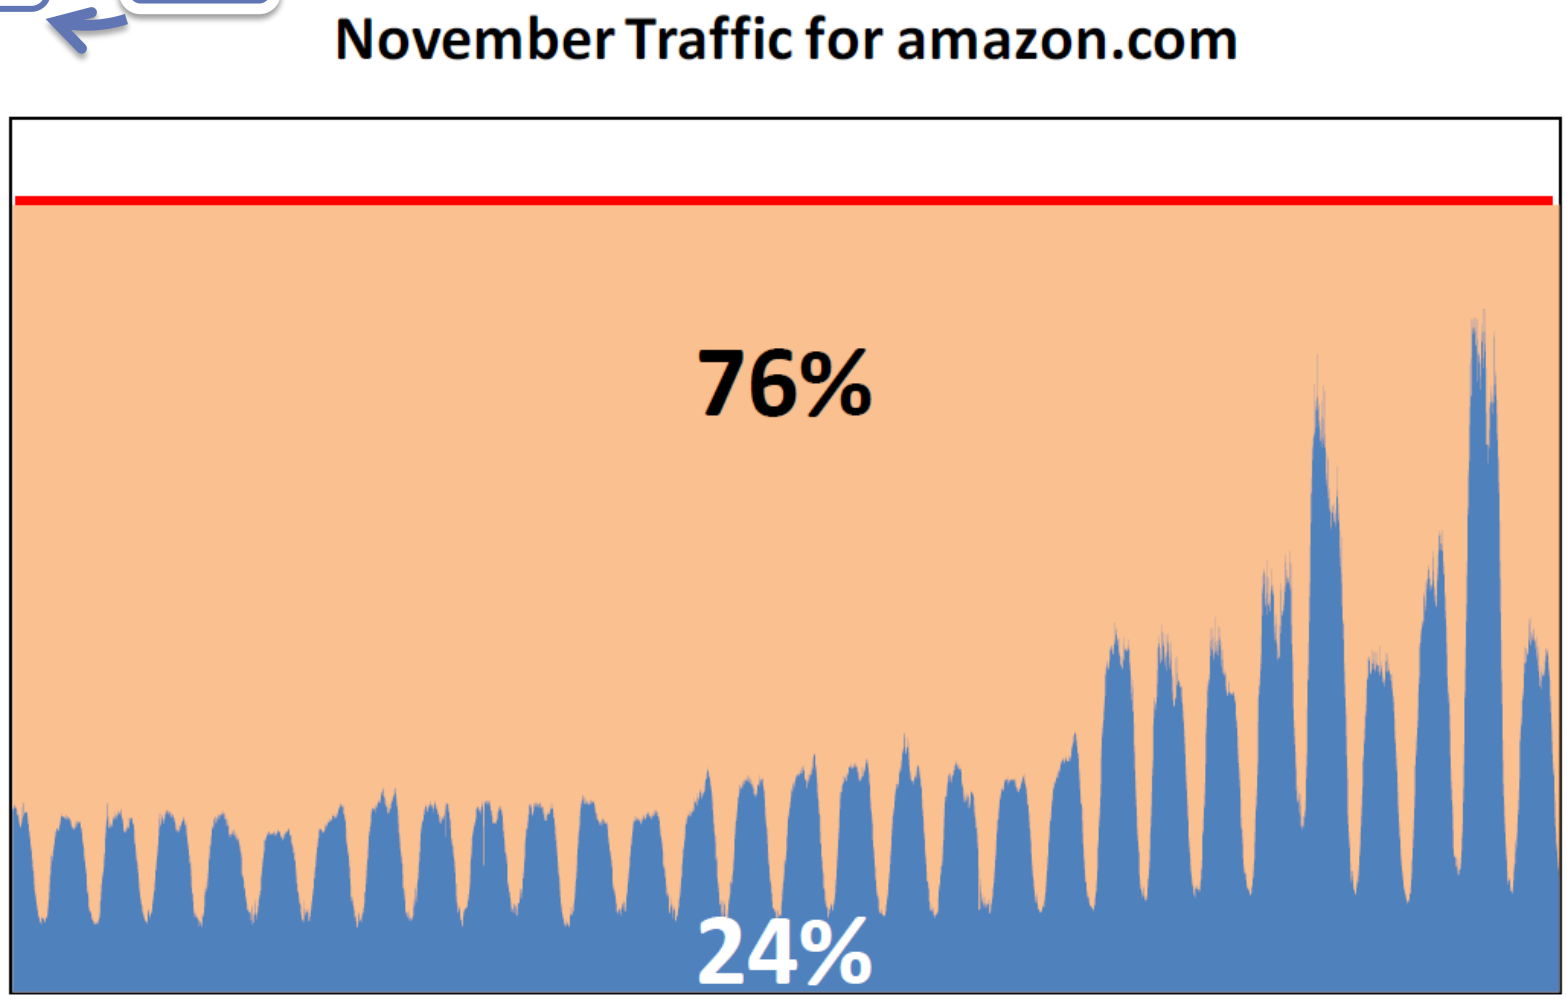
November Traffic for amazon.com
76%
24%

Might as well be Flushing $$

- Must have Active Passive Failover
- Let's Build for Peak + Buffer
- You Can't Survive if your COGS are the highest
- Give me a Break

Let's Move to the Cloud

- November 10th 2010 full migration to EC2
- Reduced spending on server capacity
- Fleet scales dynamically in increments as small as a single host
- Traffic spikes handled with ease
- Cultural change – aim for small server footprints

Business Continuity

A Cautionary Tale

Microsoft Acquires Farecast For $115M

by **Mark Hendrickson** on April 17, 2008 **50 Comments** 1 retweet Share 3

Rumors about the acquisition of **Farecast** are accurate – in a very brief blog **post** CEO **Hugh Crean** says they've been acquired by Microsoft.

SeattlePI, which first **broke** the rumor last week, **says** the price tag was $115 million. While the two companies are an understandable fit given their proximity and partnership over MSN Travel, SeattlePI reports that Farecast entertained multiple offers before accepting Microsoft's.

Farecast is an airfare pricing comparision tool

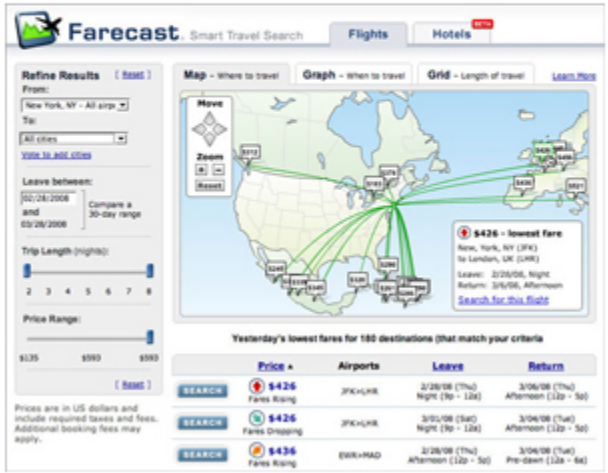

April 2008

Farecast becomes Bing Travel

No Safety Net

Service housed in a single Datacenter.

No Budget for 2nd DC Buildout.

July 2009
Disaster Strikes!
An Electrical Fire @ Fisher Plaza

TV Stations, Radio Stations,
Online Games, & Bing Travel

Bing Travel is now

2+ Datacenters

Reward

Mitigation

Risk

Microsoft has Geo-Redundancy

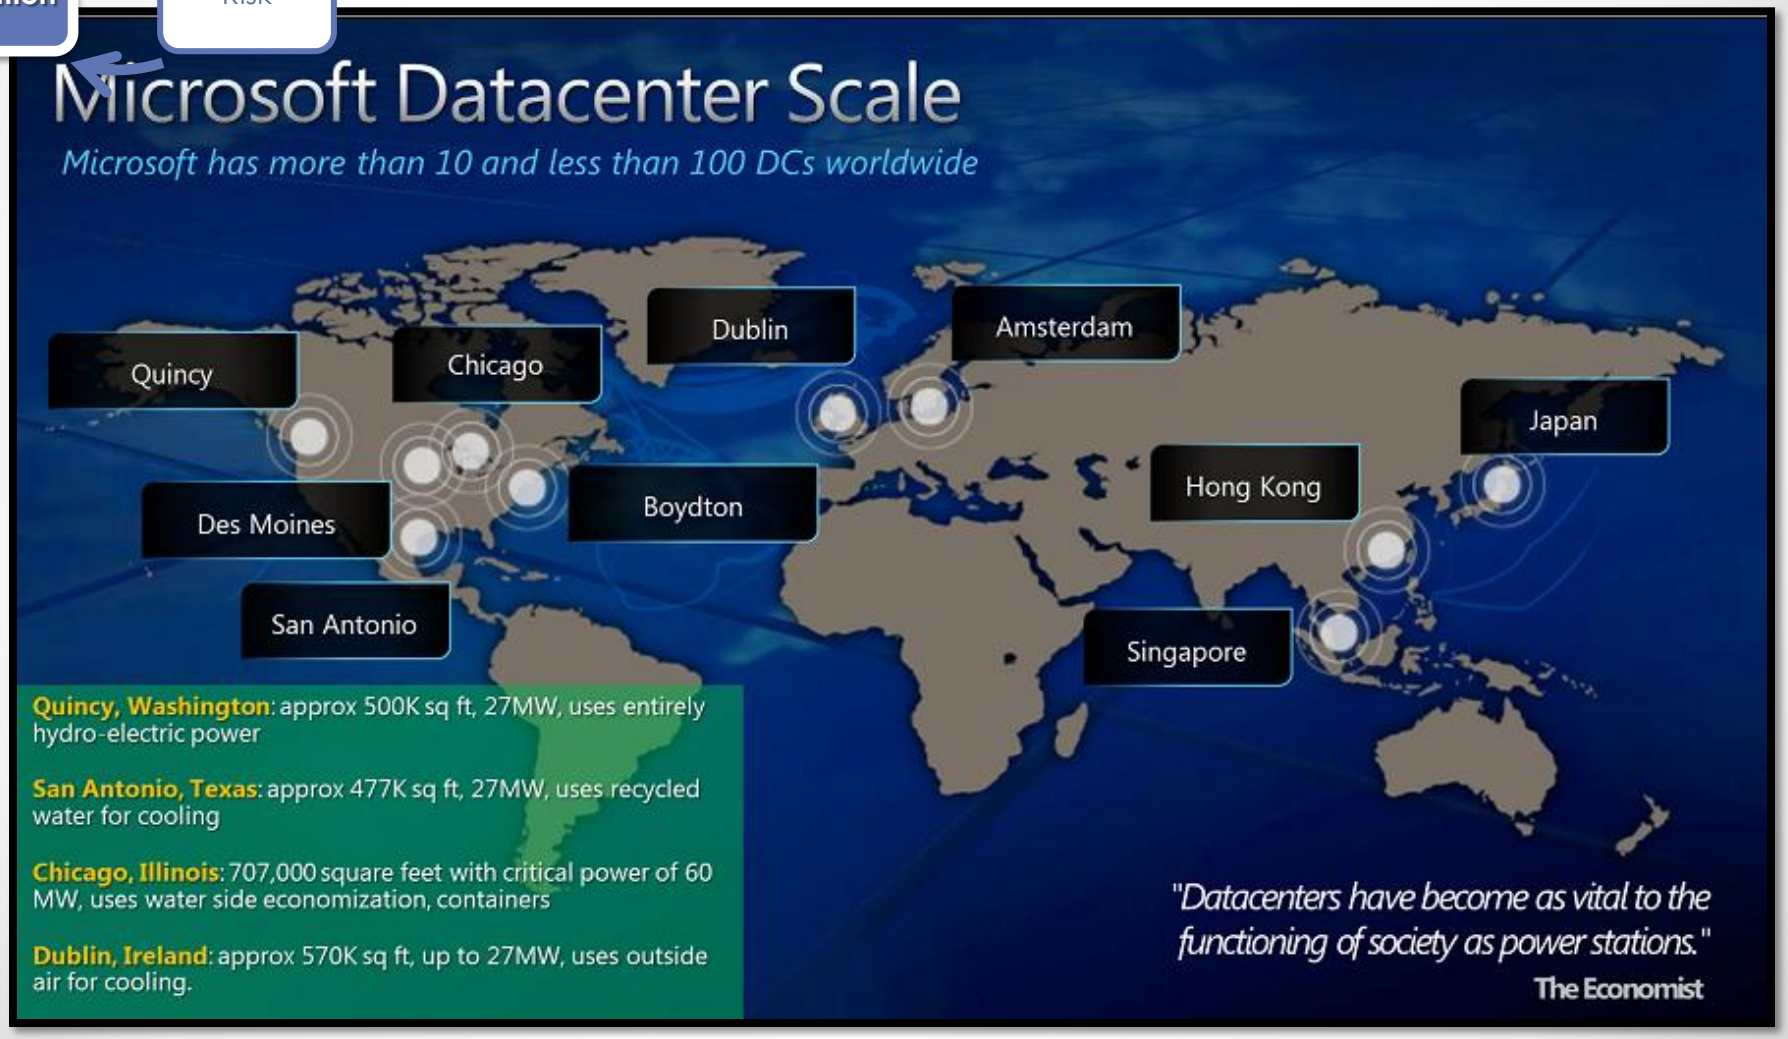

…Therefore YOU have Geo-Redundancy

…in The Cloud

- Windows Azure Traffic Manager
  - Automatically **load balance** traffic to the best data center
    - Performance
    - Failover

- Amazon S3 Storage
  - "data is **replicated** over multiple locations such that failure modes are independent of each other. The locations are chosen with great care to achieve this independence"

[Amazon geo, May 2010]

- Google Cloud Storage
  - "We **replicate** data to multiple data centers and serve an end-user's request from the nearest data center that holds a copy of the data"

[Google Cloud Storage]

[Images: http://thejoyofcode.com]

…Or Do You?

Again, you are responsible for good design

April 21, 2011 – Skynet begins it's attack against humanity

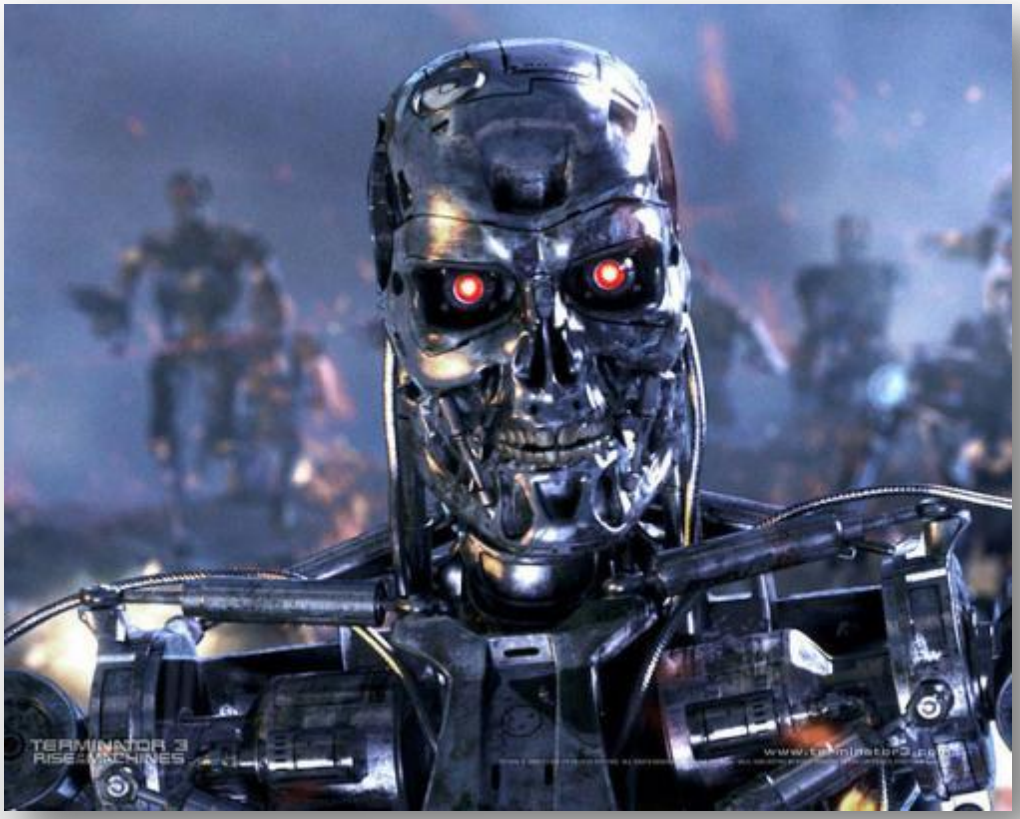

http://en.wikipedia.org/wiki/Skynet

Credit to Don MacAskill for pointing this out

…Or Do You?

Again, you are responsible for good design

April 21, 2011 – Amazon AWS EC2/RDS Outage

- Took down

reddit hootsuite foursquare Quora

- But one website had reason to be SmugMug
  - …minimally impacted, and all major services remained online during the AWS outage
- Netflix stayed up too… more later…

…Or Do You?

Again, you are responsible for good design

April 21, 2011 – Amazon AWS EC2/RDS Outage

- Took down
- Bu
  -
- Netflix stayed up too… more later…

You must…
Design for Redundancy

Don't Be This Guy

Life of our patients is at stake - I am desperately asking you to contact

Posted by: md76040303317
Posted on: Apr 22, 2011 11:20 PM

This question is **answered**. Helpful answers available: **2**. Correct answers available: **1**.

Sorry, I could not get through in any other way

We are a monitoring company and are monitoring hundreds of cardiac patients at home. We were unable to see their ECG signals since 21st of April

Could you please contact us?

How Did SmugMug Do It?

Amazon EC2 (N. California)	✓	✓	✓
Amazon EC2 (N. Virginia)	⛔	⛔	⛔

- Availability Zones (AZs)
- Failures Should Not Span AZs
  - In this case they did!
- SmugMug uses Three AZs
- Designed to fail and recover
  - Any of our instances, or any group of instances in an AZ, can be "shot in the head" [SmugMug April 2011]
- Incident Response
  - We updated our own status board, and then I tried to work around the problem…. 5 minutes [later] we were back in business

How Do You Do It?

Service [Sub-Region]
Compute [East Asia]
Compute [East US]
Compute [North Central US]
Compute [North Europe]
Compute [South Central US]
Compute [Southeast Asia]
Compute [West Europe]
Compute [West US]

- Multiple Amazon AZs or Azure Regions
- Plus Traffic Management
  - Multiple service instances costs more

- Azure LRS: Local Redundant Storage
  - Protects against common failures (disk, node, rack)
- Azure GRS: Geo-Redundant Storage
  - Protects against Data Center outage
  - Costs 23%-34% more

Choose how to spend your $$$
  - Resiliency or Response

Fault Tolerance

..or What Do You Need to Worry About When Running Your Own Data Center

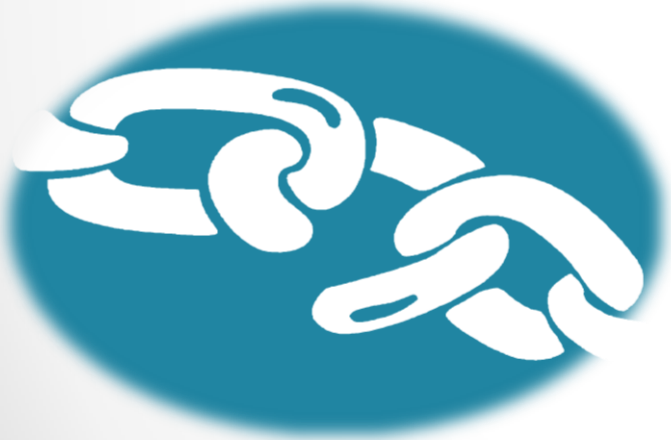

Failure is Always an Option

For Example…. NETFLIX

First Year -New Data Center

- 1 Power Distribution Unit failure (500-1000 machines)
- 1 rack-move (500-1000 machines)
- 1 network rewiring (rolling 5% of machines)
- 20 rack failures (40-80 machines)
- 8 network maintenances (~30-min connectivity losses)
- 12 router reloads
- 3 router failures
- Dozens of minor 30-second blips for DNS
- 1000 individual machine failures
- 1000s of hard drive failures

[Google Cluster, 2008]

How Does The Cloud Help?

The Cloud **is** better

- Fault-tolerant hardware and network infrastructure
- Advanced Ops personnel and processes
- State of the art: Power, Cooling, Security

amazon

The Cloud **is *not*** better

- but gives you better tools to….

…Embrace Failure

aka
design defensively

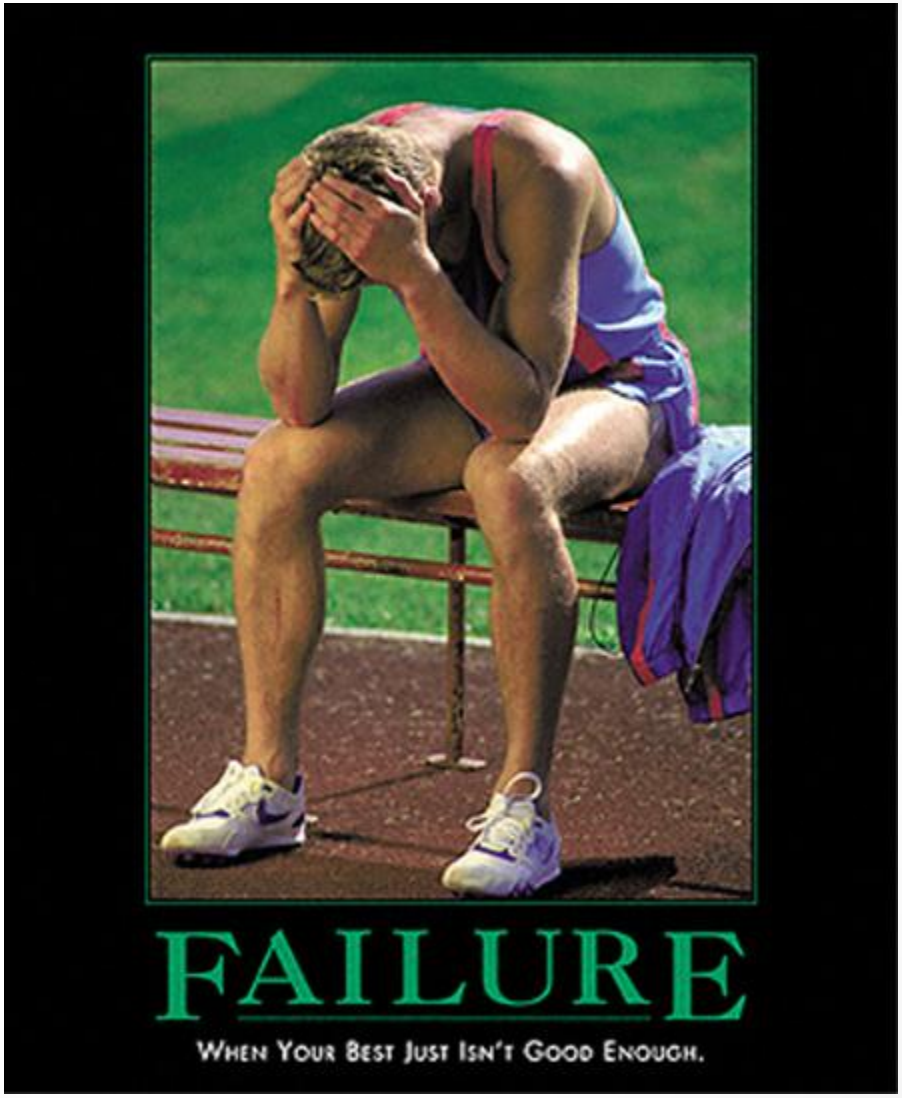

[http://despair.com]

Embrace Failure

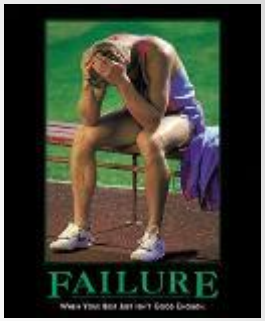

Design Defensively

- Each System has to succeed, even on its own
  - Small Stateless Services
  - Recommendation System Down? Show popular titles instead of personalized picks
- Assume host failures happen
  - Remember, "shot in the head"
  - Cloud Advantage: Re-Spawn!
- Short Timeouts and Quick Retries – Fail Fast
  - Co-tenancy can introduce variance in throughput at any level of the stack.
  - Requires Idempotent Interfaces
- Research and Test with Full Scale / Real Data
  - Cloud Advantage: Elasticity

[Netflix AWS, Dec 2010]
[Twilio AWS, Apr 2011]

Destructive Testing

[Netflix Army, July 2011]

Netflix Simian Army

- **Chaos monkey** randomly disables production instance in AWS
- **Chaos Gorilla** simulates an outage of an entire Amazon AZ
- Janitor Monkey, Security Monkey, Latency Monkey…..

Amazon Game Day

- An entire Data Center is "wiped out":

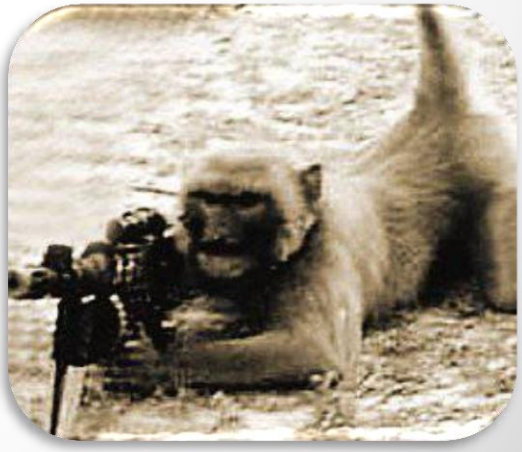


Reward
Risk
Mitigation
amazon
web services™
NETFLIX

Security

"...every cloud customer retains responsibility for assessing and understanding the value and sensitivity of the data they may choose to move to the cloud. As the owners of that information, cloud customers also remain accountable for decisions regarding the protection of that data wherever it may be stored."

[Microsoft Security, 2010]

For Example….

Amazon AMIs

Amazon Machine Image

- Create and share virtual server configurations
- Like Open Source –Give a little, Get a lot

AMI Key Vulnerability

AMI = House
SSH Key to House

June 2008 [Cloud Security 2008]

- User creates AMI
- AMI uploaded to AWS
- Other users use AMI
- Amazon Closes "Hole"

Abundant Security Problems?

June 2011

- Users Publish AMIs containing API Authentication Keys

- Amazon's or User fault?
  - User Violated Amazon Security Guideline

[IT World, 2011]

Amazon AMI Mitigation

Testing in The Cloud

Facebook is a Cloud Platform

Apps power Facebook

heroku + facebook =

Deploy and Run FB Apps [FB Heroku, 2011]

This is PaaS

Rewards:

- Supports Ruby, Node.js, Python, or PHP
- No need to setup host
- Instant Scaling

What are the Risks?

How do We Test it?

- Does it work?
- Is it stable?
- Users getting a Good Experience?

These risks are not cloud specific.

But the mitigation is….

Imaginary Friends

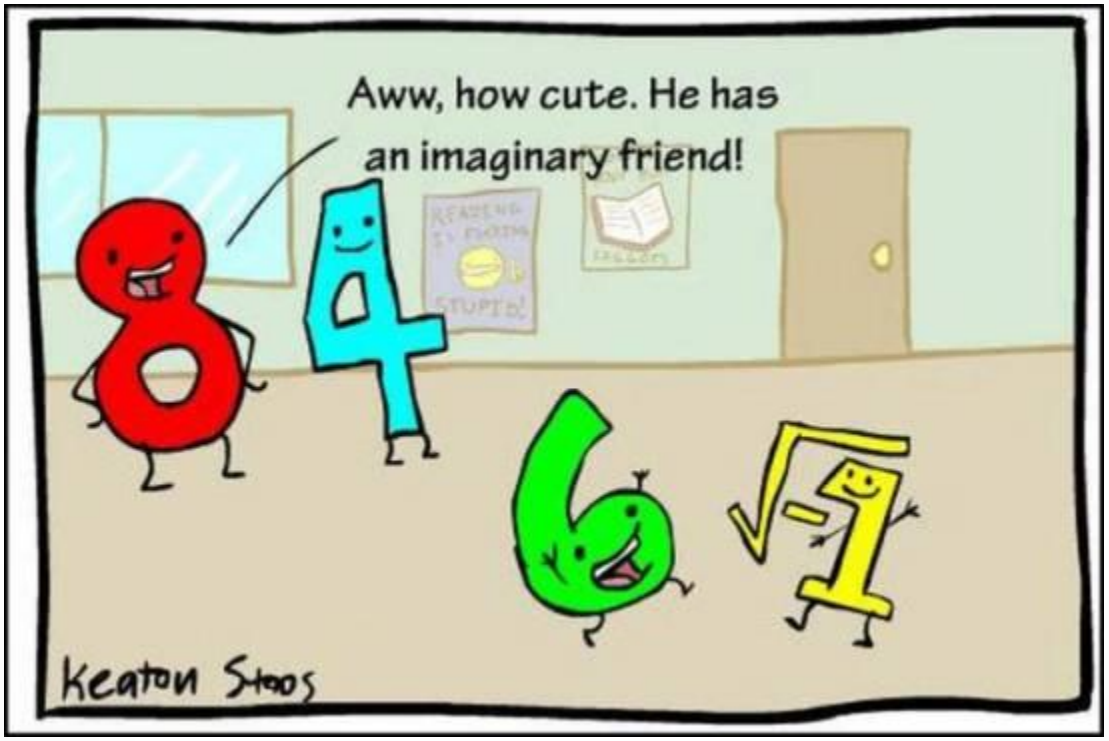

Facebook Imaginary Friends

…they call them **Test Users**

- Invisible user accounts
- Not visible by others; can only be friends with other Test Users
- Experience your app as a regular user

Power of the Cloud

- Automated:
  - Programmatic interface
  - Web UI
- Create up to 500 of them

[FB Test, 2011]

Control 1 Million Users

Control 1 Million Users

- Uses Cloud IaaS Providers:
  - GoGrid, Windows Azure, Amazon EC2
- Generate high scale load from geo-dispersed origins

- 1 million **concurrent** virtual users
  - Plus Live Traffic
- 6 gigabits per second
- 6 terabytes of data transferred per hour
- Over 77,000 hits per second Plus Live Traffic
- 800 Amazon EC2 instances / 3200 cloud computing cores

[SOASTA, 2010]

Virtual Sandbox

- Production Environment
- Staging Environment
- **Dev and Testing Environment**

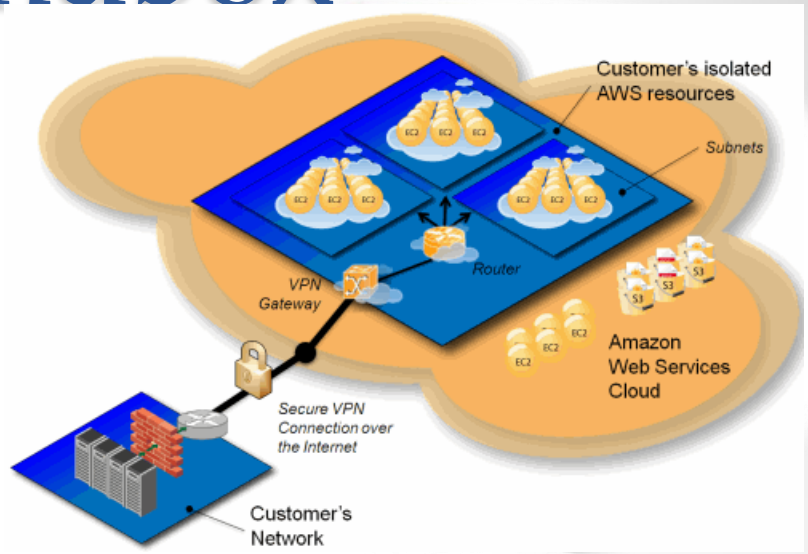

Can you have it all in one big Cloud?

- Amazon Virtual Private Cloud (Amazon VPC)
- Provision a private, isolated section of AWS
- IP addresses, subnets, routing tables
- Even Sandbox for Non-Cloud services

And remember the power of zero!

Netflix "Canary" Deployment

1B API requests per day

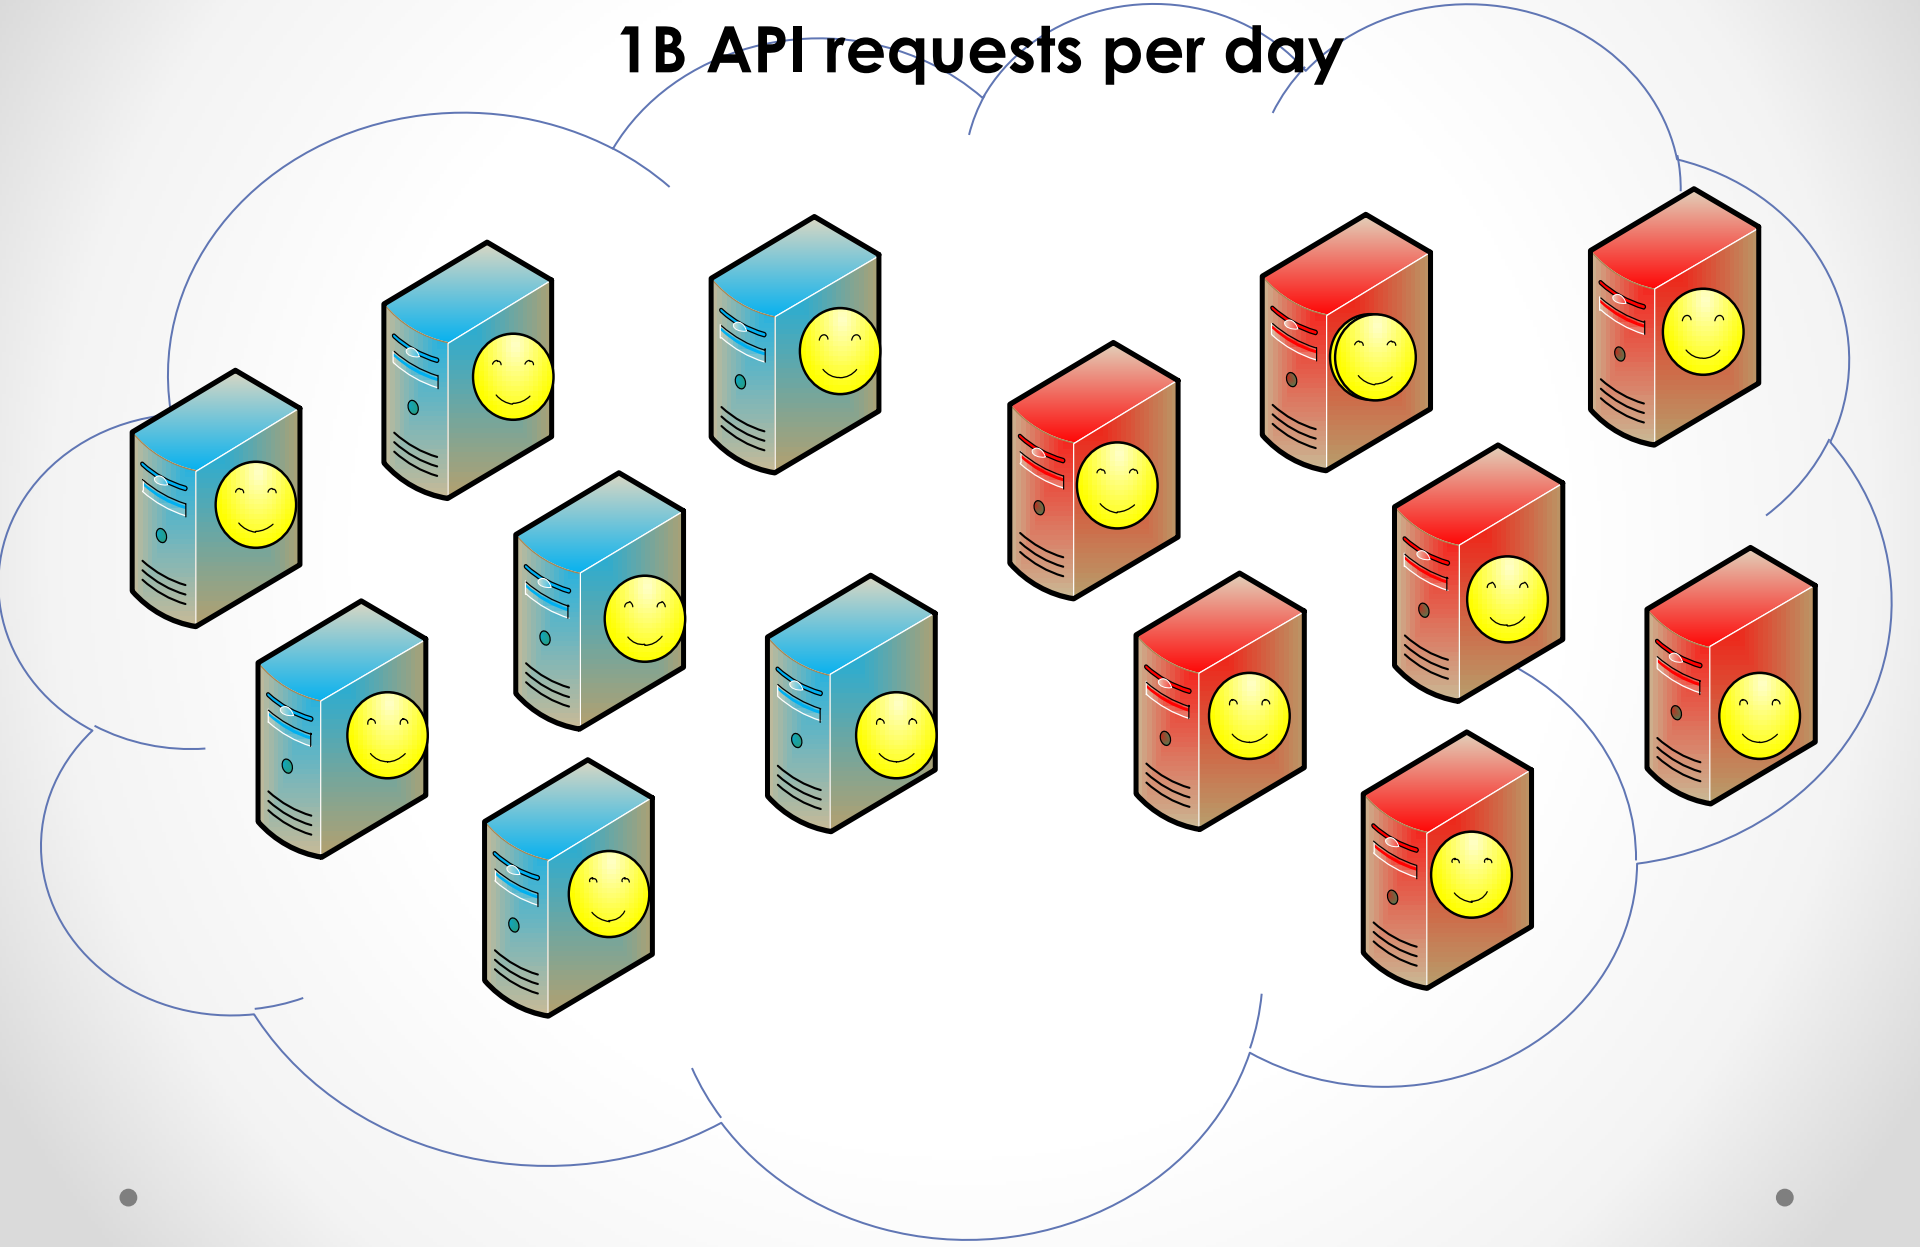

Test Oriented Architecture

Even Cloud Services need Testing

Ken's Services Theorem

- Services are like Ogres
- Ogres are like Onions
- Onions have Layers
- Therefore services have Layers

The Problem is

- The layers of a service spin at different rates
- Movement toward continuous deployment

Code Churn Example 1

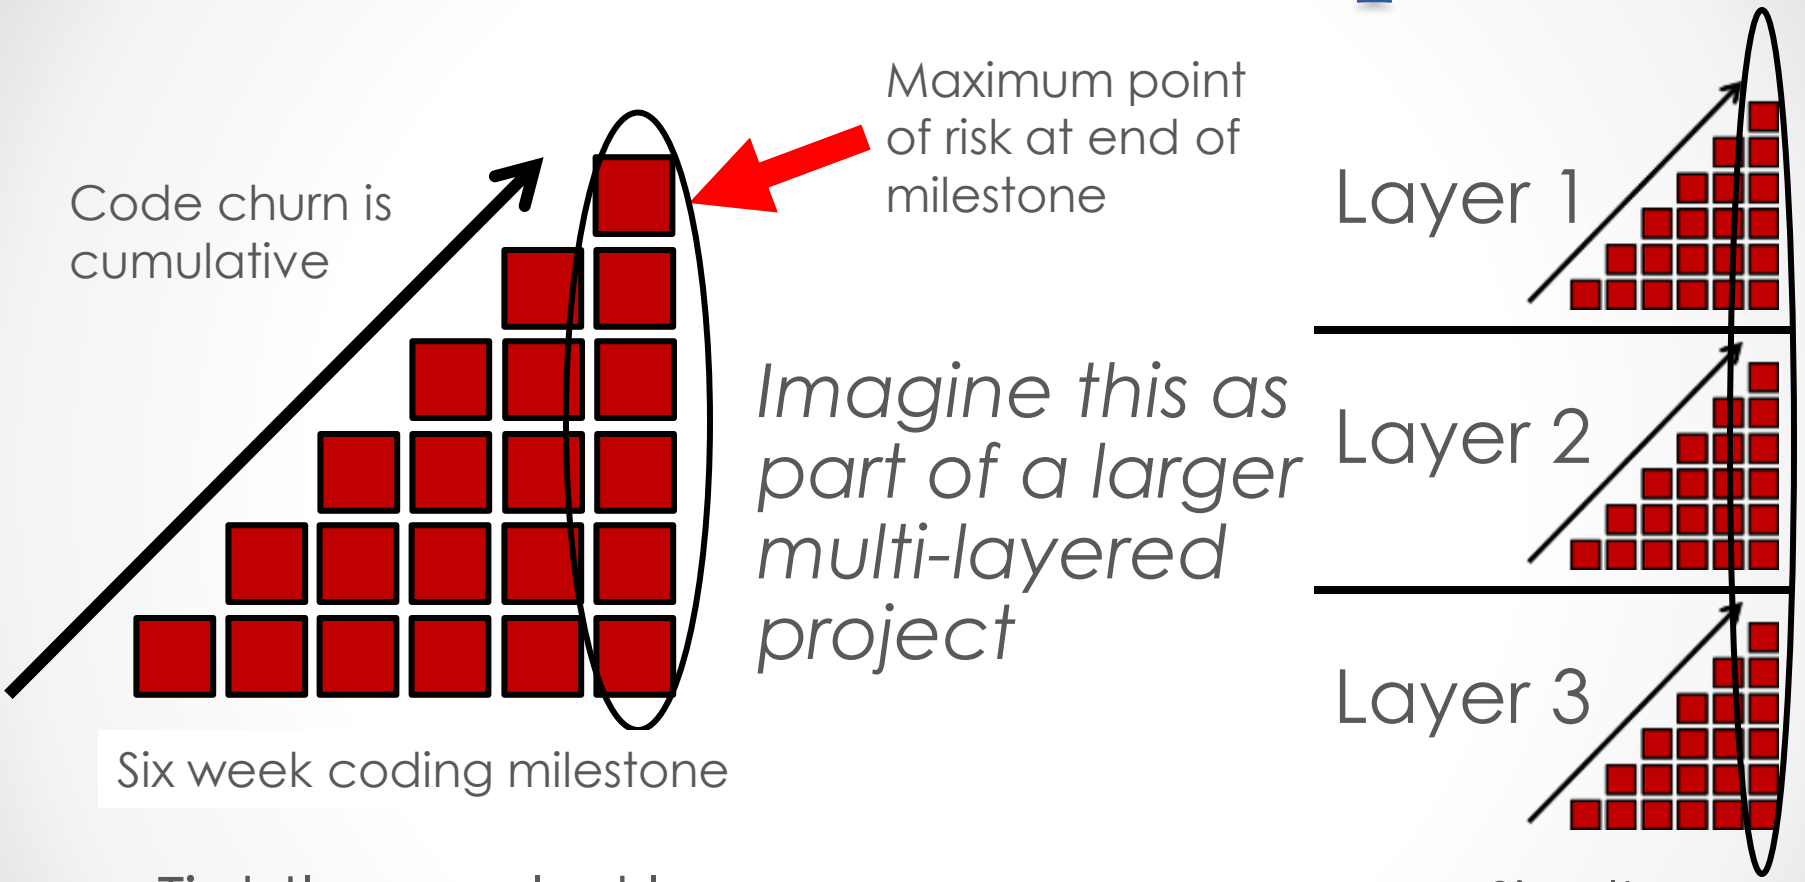

- Tightly coupled layers
- Long stabilization phase
- Complicated end-to-end integration

Code Churn Example 2 (CD)

Rapid release cadence (weekly or daily)

- Risk per release decreases because of more incremental change
- Change builds over time in production
- Next release is always the most risky

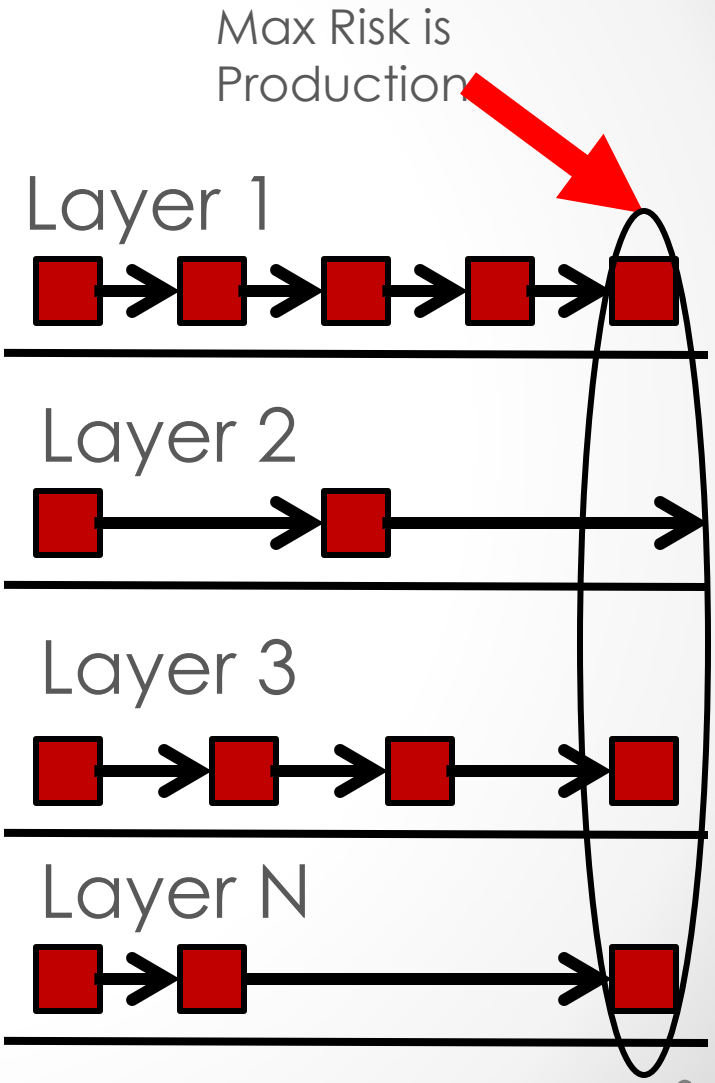

Practical TOA

- More Loose Coupling across stack
  - Splitting can be a good thing
  - Your service in the Cloud

- More Self Service Deployments
  - Automated roll forward
  - Rollback triggered by live site monitors
  - Canary deployment zones

Practical TOA

- Automated Tests and Monitors are the same thing

Heavy Test Automation

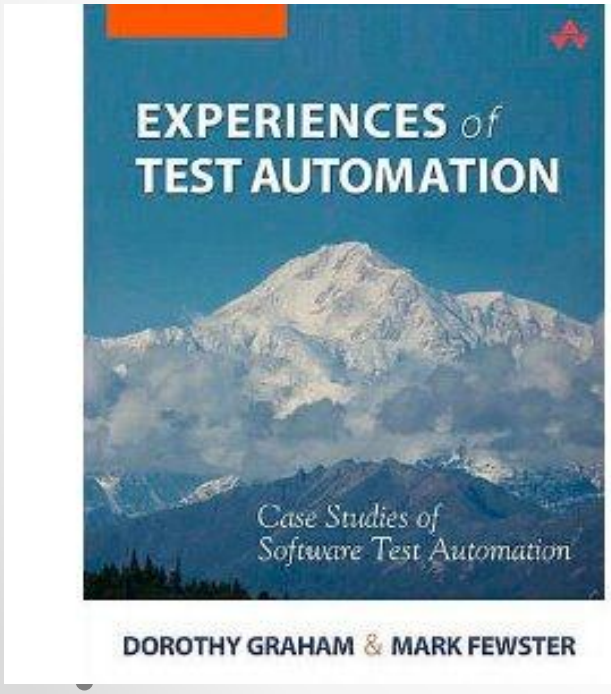

=

Big Live Service Monitors

Practical TOA

- Ship Test Hooks into production
  - Runtime Flags to access test path
  - Isolated Data Centers and Hosts
  - Runtime routing of traffic from v-Current to v-Next
- Rich Telemetry
  - Your services telemetry
  - Runtime flags for richer debug telemetry
  - Fix the bugs users are seeing

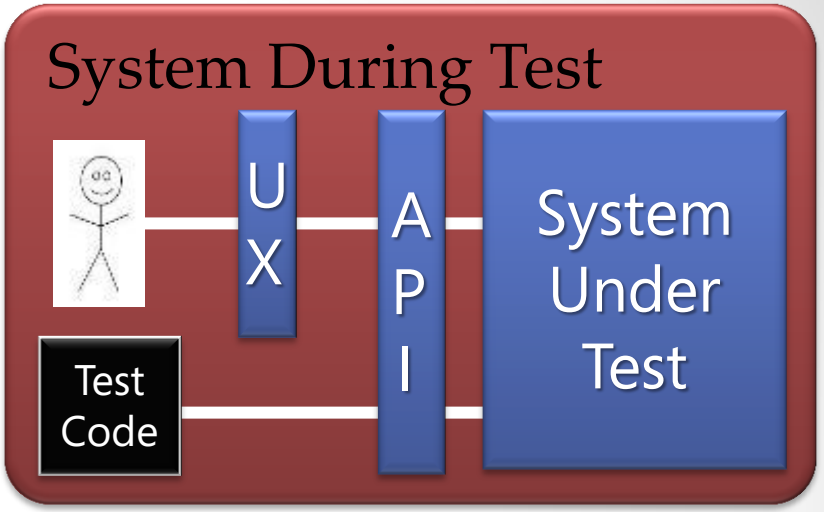

From *Alan Myrvold* **"Patterns of Testability"**

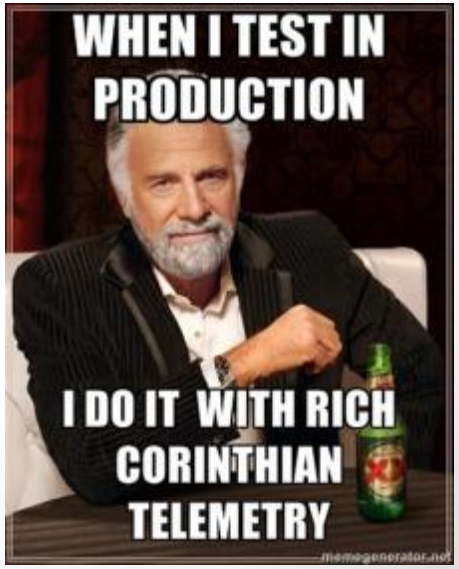

Summary

- About Clouds
- Cloud Rewards
- Getting Into The Cloud
- 5 Amazing Cloud Case Studies
  - Rewards, Risks & Mitigations
- Testing in The Cloud

The latest version of this slide deck can be found at:
http://www.setheliot.com/blog/bsc-west-2012/

References

[Amazon geo, May 2010]	Expanding the Cloud - Amazon S3 Reduced Redundancy Storage, Werner Vogels May 2010; http://www.allthingsdistributed.com/2010/05/amazon_s3_reduced_redundancy_storage.html
[Amazon Growth, 2011]	Amazon S3 - 566 Billion Objects, 370,000 Requests/Second, and Hiring! Oct 4, 2011 http://aws.typepad.com/aws/2011/10/amazon-s3-566-billion-objects-370000-requestssecond-and-hiring.html
[AWS Whitepaper]	Migrating your Existing Applications to the AWS Cloud (with 3 example scenarios) Oct 2010; http://d36cz9buwru1tt.cloudfront.net/CloudMigration-main.pdf
[Berkeley 2009]	Above the Clouds: A Berkeley View of Cloud Computing; Feb 2009; http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf
[CCJ, 2011]	http://cloudcomputing.sys-con.com/node/1662284 , Feb 2011
[Cloud Security 2008]	Is Your Amazon Machine Image Vulnerable to SSH Spoofing Attacks?, July 2008; http://cloudsecurity.org/tags/ssh.html
[Cloud SLAs]	http://www.microsoft.com/windowsazure/sla/; http://aws.amazon.com/ec2-sla/; http://www.rackspace.com/cloud/legal/sla/; http://www.google.com/apps/intl/en/terms/sla.html
[Cloud Tweaks, 2011]	http://www.cloudtweaks.com/2011/08/3-handy-cloud-computing-price-comparison-engines/, August 2011
[Deschamps 2012]	Experiences of Test Automation; Dorothy Graham; Jan 2012; ISBN 0321754069; Chapter: "Moving to the Cloud: The Evolution of TiP, Continuous Regression Testing in Production"; Ken Johnston, Felix Deschamps
[FB Heroku, 2011]	Facebook and Heroku; http://blog.heroku.com/archives/2011/9/15/facebook/, Sept 15 2011 ; https://devcenter.heroku.com/articles/facebook
[FB Test, 2011]	Making it easier to create and manage Test Users ; http://developers.facebook.com/blog/post/527/ , July 27 2011
[FISMA, 2002]	http://en.wikipedia.org/wiki/Federal_Information_Security_Management_Act_of_2002
[Forbes, 2011]	http://www.forbes.com/sites/kevinjackson/2011/04/19/cloud-to-command-90-of-microsofts-rd-budget/ , April 2011
[Google Cloud Storage]	http://googledevelopers.blogspot.com/search/label/google%20storage
[Google Cluster, 2008]	Jeff Dean, Google IO Conference 2008, via Stephen Shankland, CNET http://news.cnet.com/8301-10784_3-9955184-7.html

References

[Hinchcliffe, 2009]	http://www.zdnet.com/blog/hinchcliffe/cloud-computing-and-the-return-of-the-platform-wars/303 , March 2009
[IT World, 2011]	Amazon's cloud is full of holes, June 2011; http://www.itworld.com/security/175927/researchers-aws-users-are-leaving-security-holes
[Jenkins, 2011]	Velocity 2011: Jon Jenkins, "Velocity Culture" , June 2011; http://www.youtube.com/watch?v=dxk8b9rSKOo
[MAP Toolkit]	Microsoft Assessment and Planning (MAP) Toolkit for Windows Azure Platform ; http://technet.microsoft.com/en-us/solutionaccelerators/gg581074
[Microsoft DC, 2011]	Microsoft GFS Datacenter Tour (4:53); http://www.youtube.com/watch?v=hOxA1I1pQIw
[Microsoft Jobs, March 2012]	http://www.microsoft.com/en-us/news/features/2012/mar12/03-05CloudComputingJobs.aspx
[Microsoft Security, 2010]	Information Security Management System for Microsoft's Cloud Infrastructure, http://www.globalfoundationservices.com/security/documents/InformationSecurityMangSysforMSCloudInfrastructure.pdf November 2010
[Netflix Army, July 2011]	The Netflix Simian Army; July 2011; http://techblog.netflix.com/2011/07/netflix-simian-army.html
[Netflix Autoscaling, 2012]	http://techblog.netflix.com/2012/01/auto-scaling-in-amazon-cloud.html
[Netflix AWS, Dec 2010]	5 Lessons We've Learned Using AWS , Dec 2010; http://techblog.netflix.com/2010/12/5-lessons-weve-learned-using-aws.html
[R&M, 2010]	http://www.researchandmarkets.com/reportinfo.asp?cat_id=0&report_id=1395650 , Oct 2010
[SmugMug April 2011]	How SmugMug survived the Amazonpocalypse, April 2011; http://don.blogs.smugmug.com/2011/04/24/how-smugmug-survived-the-amazonpocalypse/
[SOASTA, 2010]	How MySpace Tested Their Live Site with 1 Million Concurrent Users; http://highscalability.com/blog/2010/3/4/how-myspace-tested-their-live-site-with-1-million-concurrent.html , March 4 2010
[Staten, 2010]	Could Cloud Computing Get Any More Confusing?; http://blogs.forrester.com/james_staten/10-05-20-could_cloud_computing_get_any_more_confusing James Staten, Forrester Research; May 20, 2010
[Twilio AWS, Apr 2011]	Why Twilio Wasn't Affected by Today's AWS Issues, April 2011; http://www.twilio.com/engineering/2011/04/22/why-twilio-wasnt-affected-by-todays-aws-issues/

Thank You

Session BW7

Leaping into "The Cloud": Rewards, Risks, and Mitigations

Ken Johnston, Seth Eliot

Thank you for attending this session.
Please fill out an evaluation form.

Azure Status

Status History

We maintain the history of the health status for each service for the past five weeks in the form of running logs. This history is shown in the table below. Mouse over a status icon to see a detailed incident report and click on the arrow icon at the top of the table to move back and forth through the weeks.

Service Windows Azure Service M

Service [Sub-Region]	May 24	May 23	ay 18
Windows Azure Service Management [Worldwide]			

Page Last Updated: 7 Jun 2012 8:21 PM UTC

Normal service availability | Performance degradation | Serv

[RESOLVED] Service Management Partial Outage

17-May-12

10:34 PM UTC We are experiencing an issue with Windows Azure service management. We are actively investigating this issue and working to resolve it as soon as possible. Further updates will be published to keep you apprised of the situation. We apologize for any inconvenience this causes our