



# ALMFORUMSEATTLE

April 1-3, 2014 • Washington State Convention Center

*Sponsored by*



Microsoft



Scrum.org

Improving the Profession of Software Development



# Your Path To Data-Driven Quality

Seth Eliot

Principal Knowledge Engineer, Test Excellence



# About Seth



Digital  
Media  
Services

A/B Testing  
of Services

Petabytes  
Processed

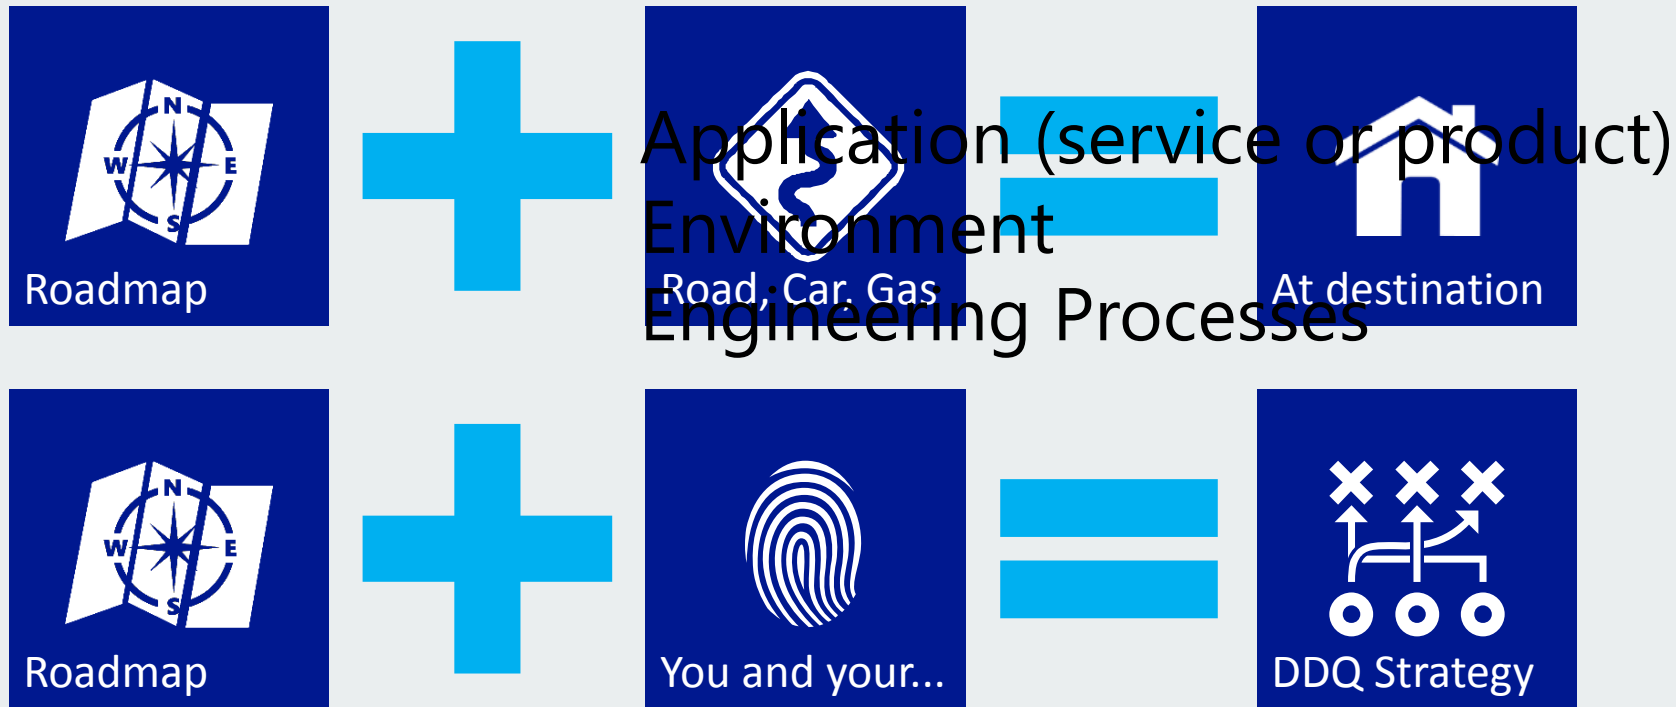
Services and  
Cloud and  
Data-Driven Quality and...





# In This Session, You'll Learn...

...how to create your "roadmap" to DDQ



# Data-Driven Quality

What is it? Why is it important?

# Data, where have we been?

## The HiPPO



## Engineering Data

Test pass/fail results

Bug counts

Delivery cadence

Code coverage

Code Churn

Highest Paid  
Person's Opinion

# Data, where are we going?

Engineering data

Test results → Scoring engines using Bayesian analysis

Production-quality data

You may have heard of TiP


Not as difficult as you might think

Lots of solutions for lots of application types

Real Users



Production Environments





# The roadmap



Determine your questions

Design for production-quality data

Select your data sources

Use the right data tools

Get answers to your questions

Learn new questions

Repeat



Performance

Availability

“Big Data” is hot, let’s start with that

Usage

?

GQM - 1994

Start at the  
beginning



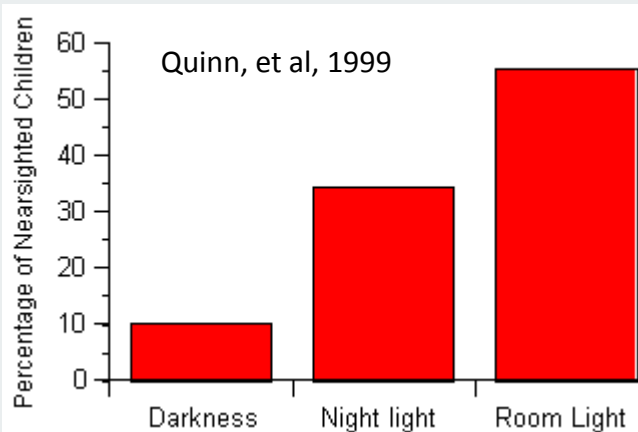
Determine your questions

Prioritize answers you seek

# Why not just get data and look for answers?



Does sunscreen increase chance of drowning?



Do night-lights cause near-sightedness in children?

# What questions does EXO ask?



Microsoft Exchange Online Dedicated Plans Version  
Service Level Agreement (SLA)

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

“Downtime” is defined as any period of time when users are unable to send or receive email via all supported mailbox access

# Is it available?

Often a Pri 0

Is the application  
(product, service)  
there for the user?



“Dialtone”



Critical work  
stoppage for  
user

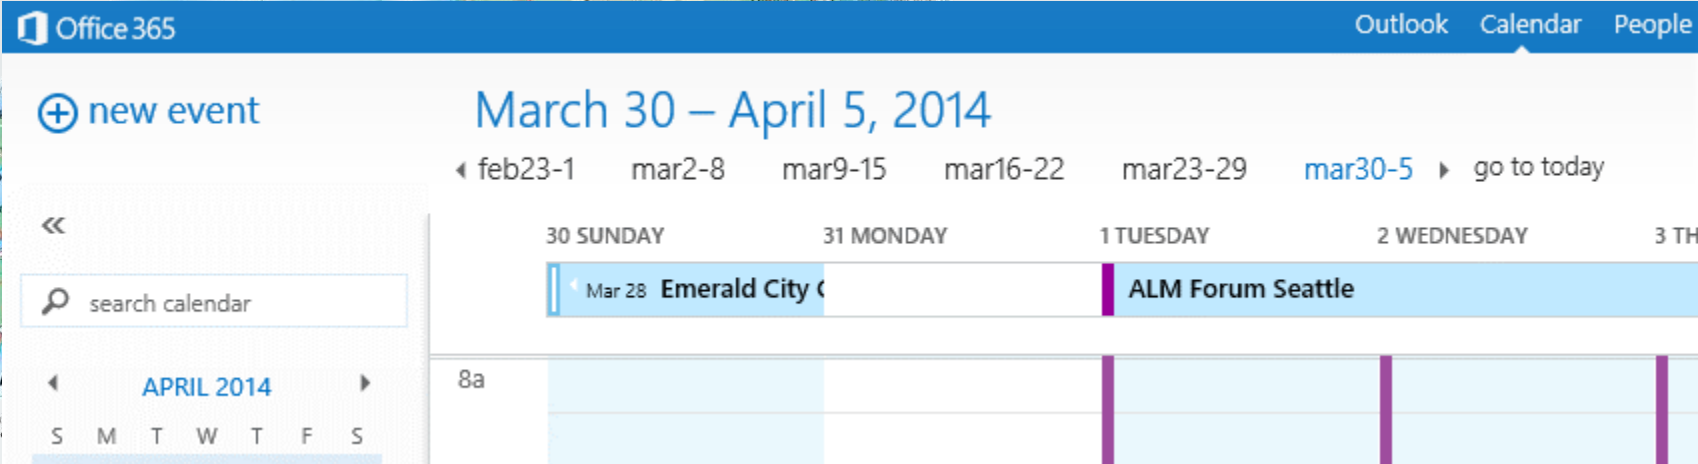
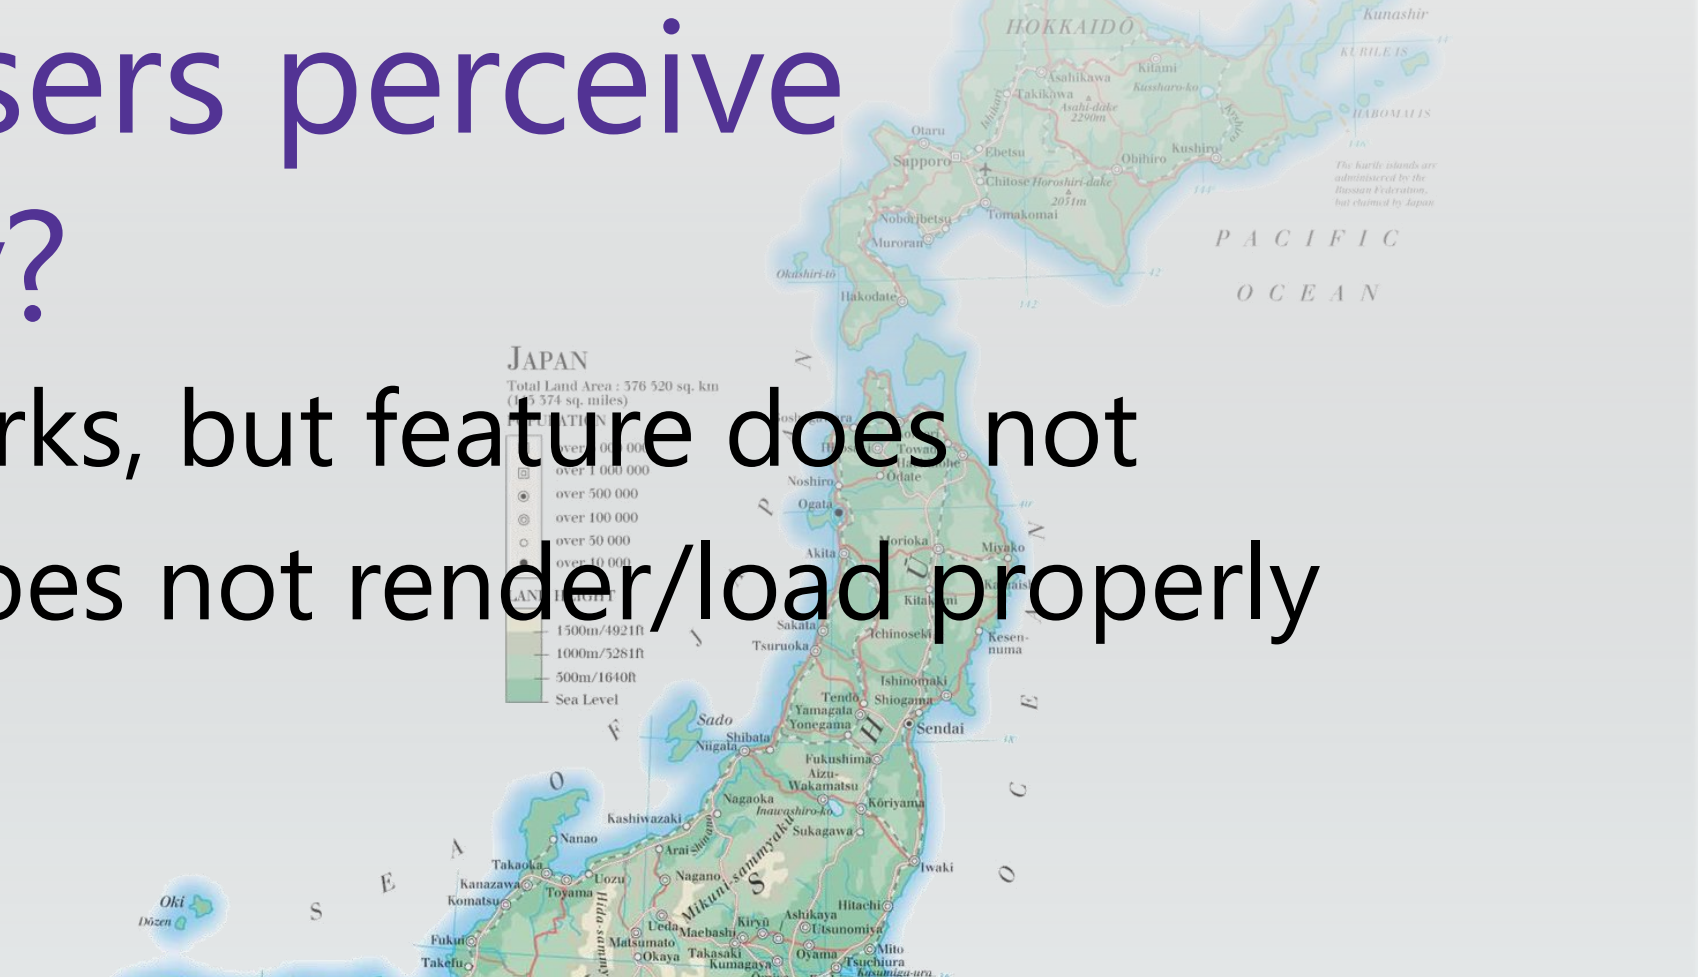




# How do users perceive availability?


Application works, but feature does not  
Occasionally does not render/load properly

違った言語




# How is performance?

## Huge Impact



## Power of Production Data



Real users



Multiple environments



End to end



Scale & geo-diversity

# What do users do?



## Customer Experience Improvement Program (CEIP)

### Customer Experience Improvement Program



#### Do you want to participate in the Windows Customer Experience Improvement Program?

The program helps Microsoft improve Windows. Without interrupting you, it collects information about your computer hardware and how you use Windows. The program also periodically downloads a file to collect information about problems you might have with Windows. The information collected is not used to identify or contact you.

[Read the privacy statement online](#)

- ☒ Yes, I want to participate in the program.  
☐ No, I don't want to participate in the program.

Save Changes

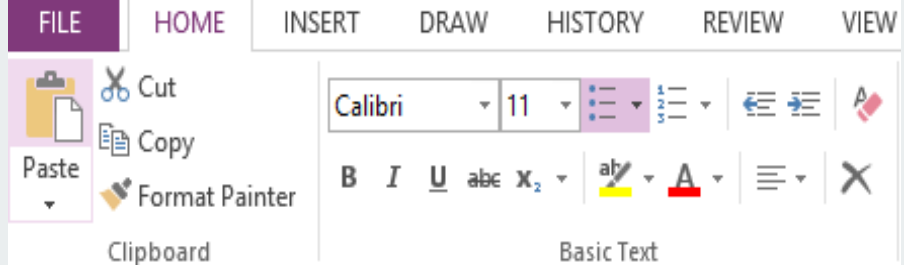
Cancel



This setting is managed by your system administrator. [Why can't I change some settings?](#)



## OneNote




# Your turn - Determine your questions



What type of answer are you looking for?

Availability, performance, usage?

Prioritize


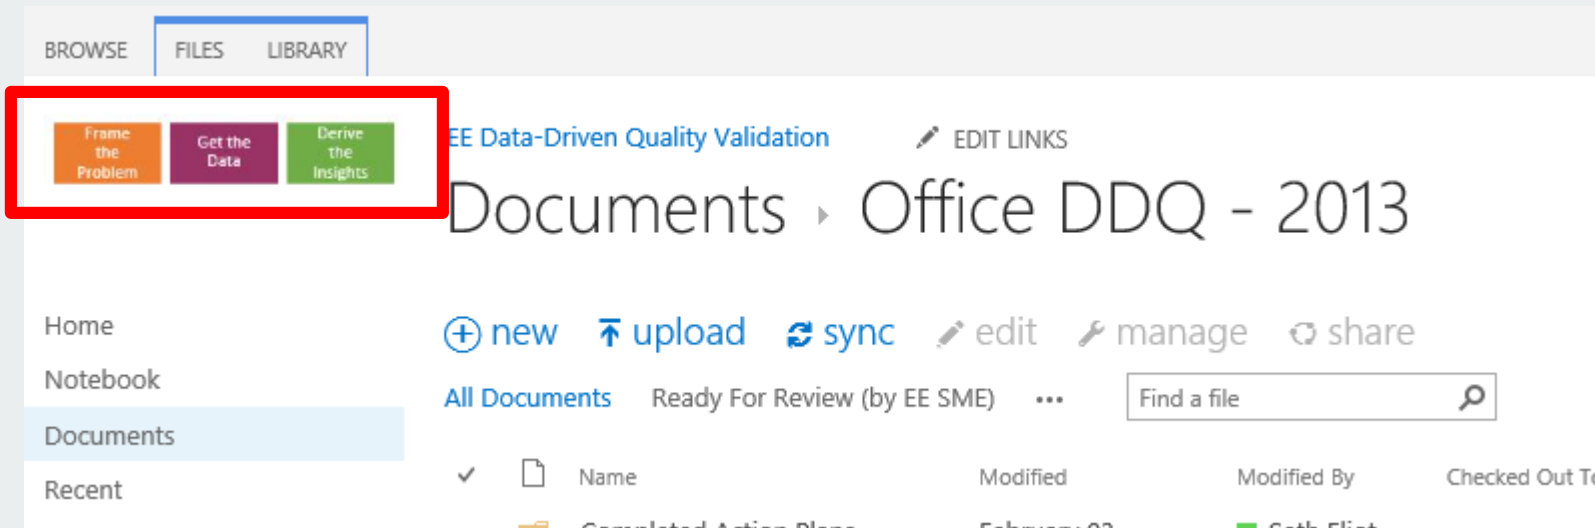


When does availability NOT come first?

# Your turn - Determine your questions



What are the key scenarios for the type you selected?



What is high priority?



# Design for production-quality data

Get data from or near production



# Two types of data to acquire


Active: synthetic

Passive: real (RUM)

Active data in prod

For services only?

Client: is the service there?



# Staged data acquisition mitigates risk


## Service



## Product (client, on-prem server)



# Staged data acquisition - Netflix



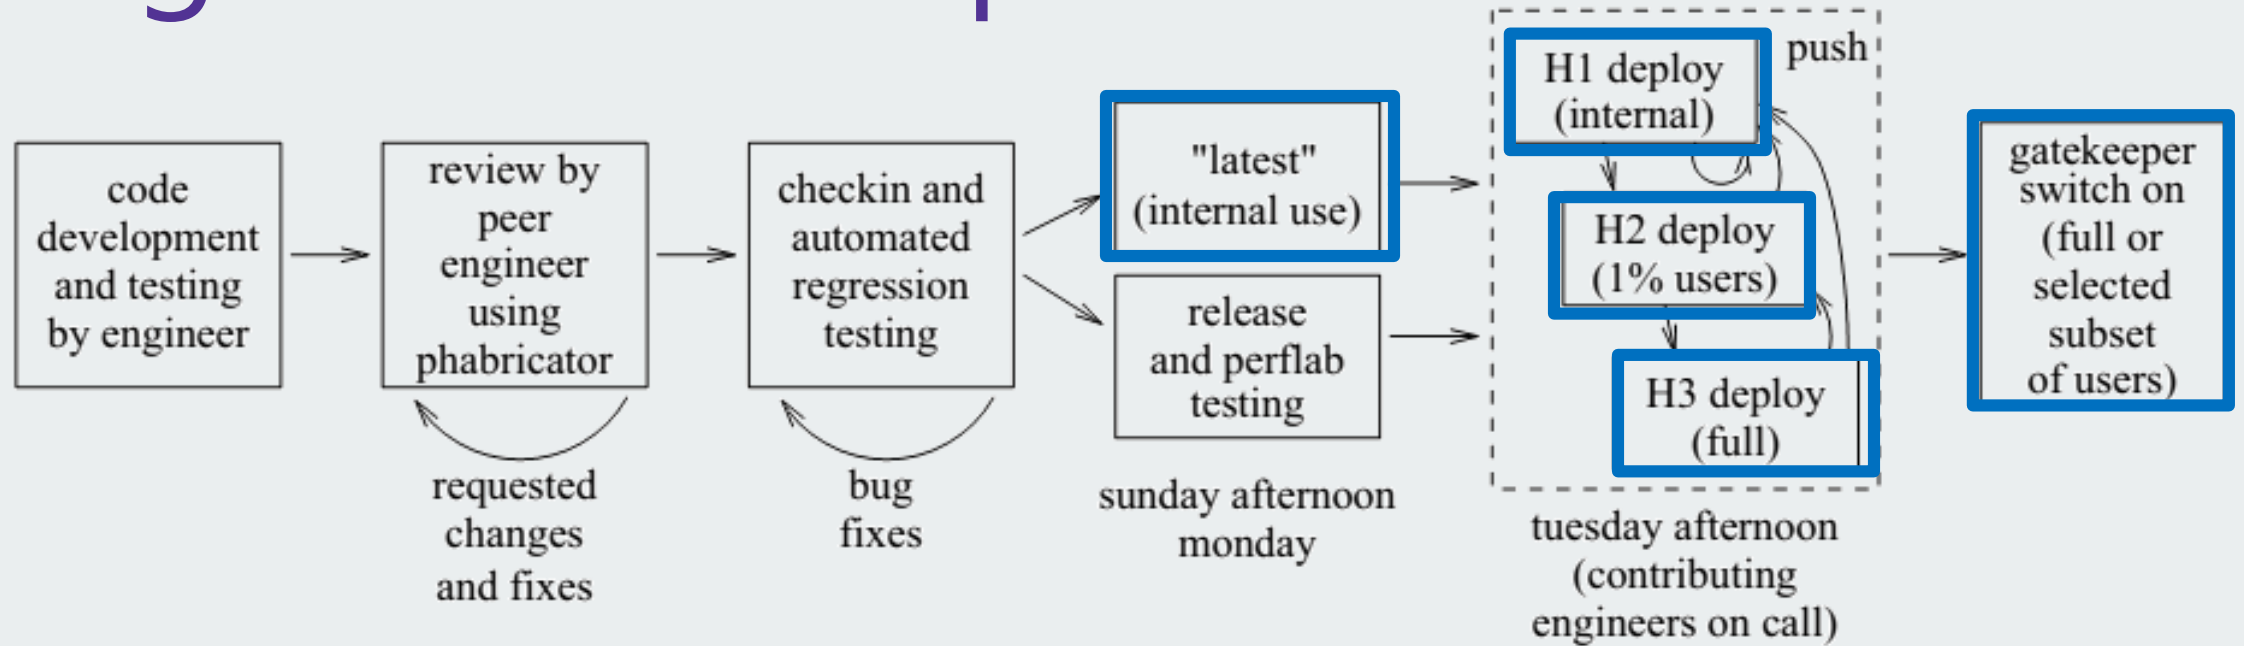
**NETFLIX**

1B API  
requests  
per day



Canary Deployment

# Staged Data Acquisition - Facebook



facebook

Dogfood


In prod, no users (except internal ones)

Some servers in Production

World-wide deployment

Feature light-up

# Speed of deployment



Usually easy for services

Client apps may have a deployment capability

Or may make use of feature light-up

EaaS - Everything is now connected and thus updatable

# Staged data acquisition - Outlook

Filtering and aggregation at client

Be kind to the client

Pipeline to collect and process data

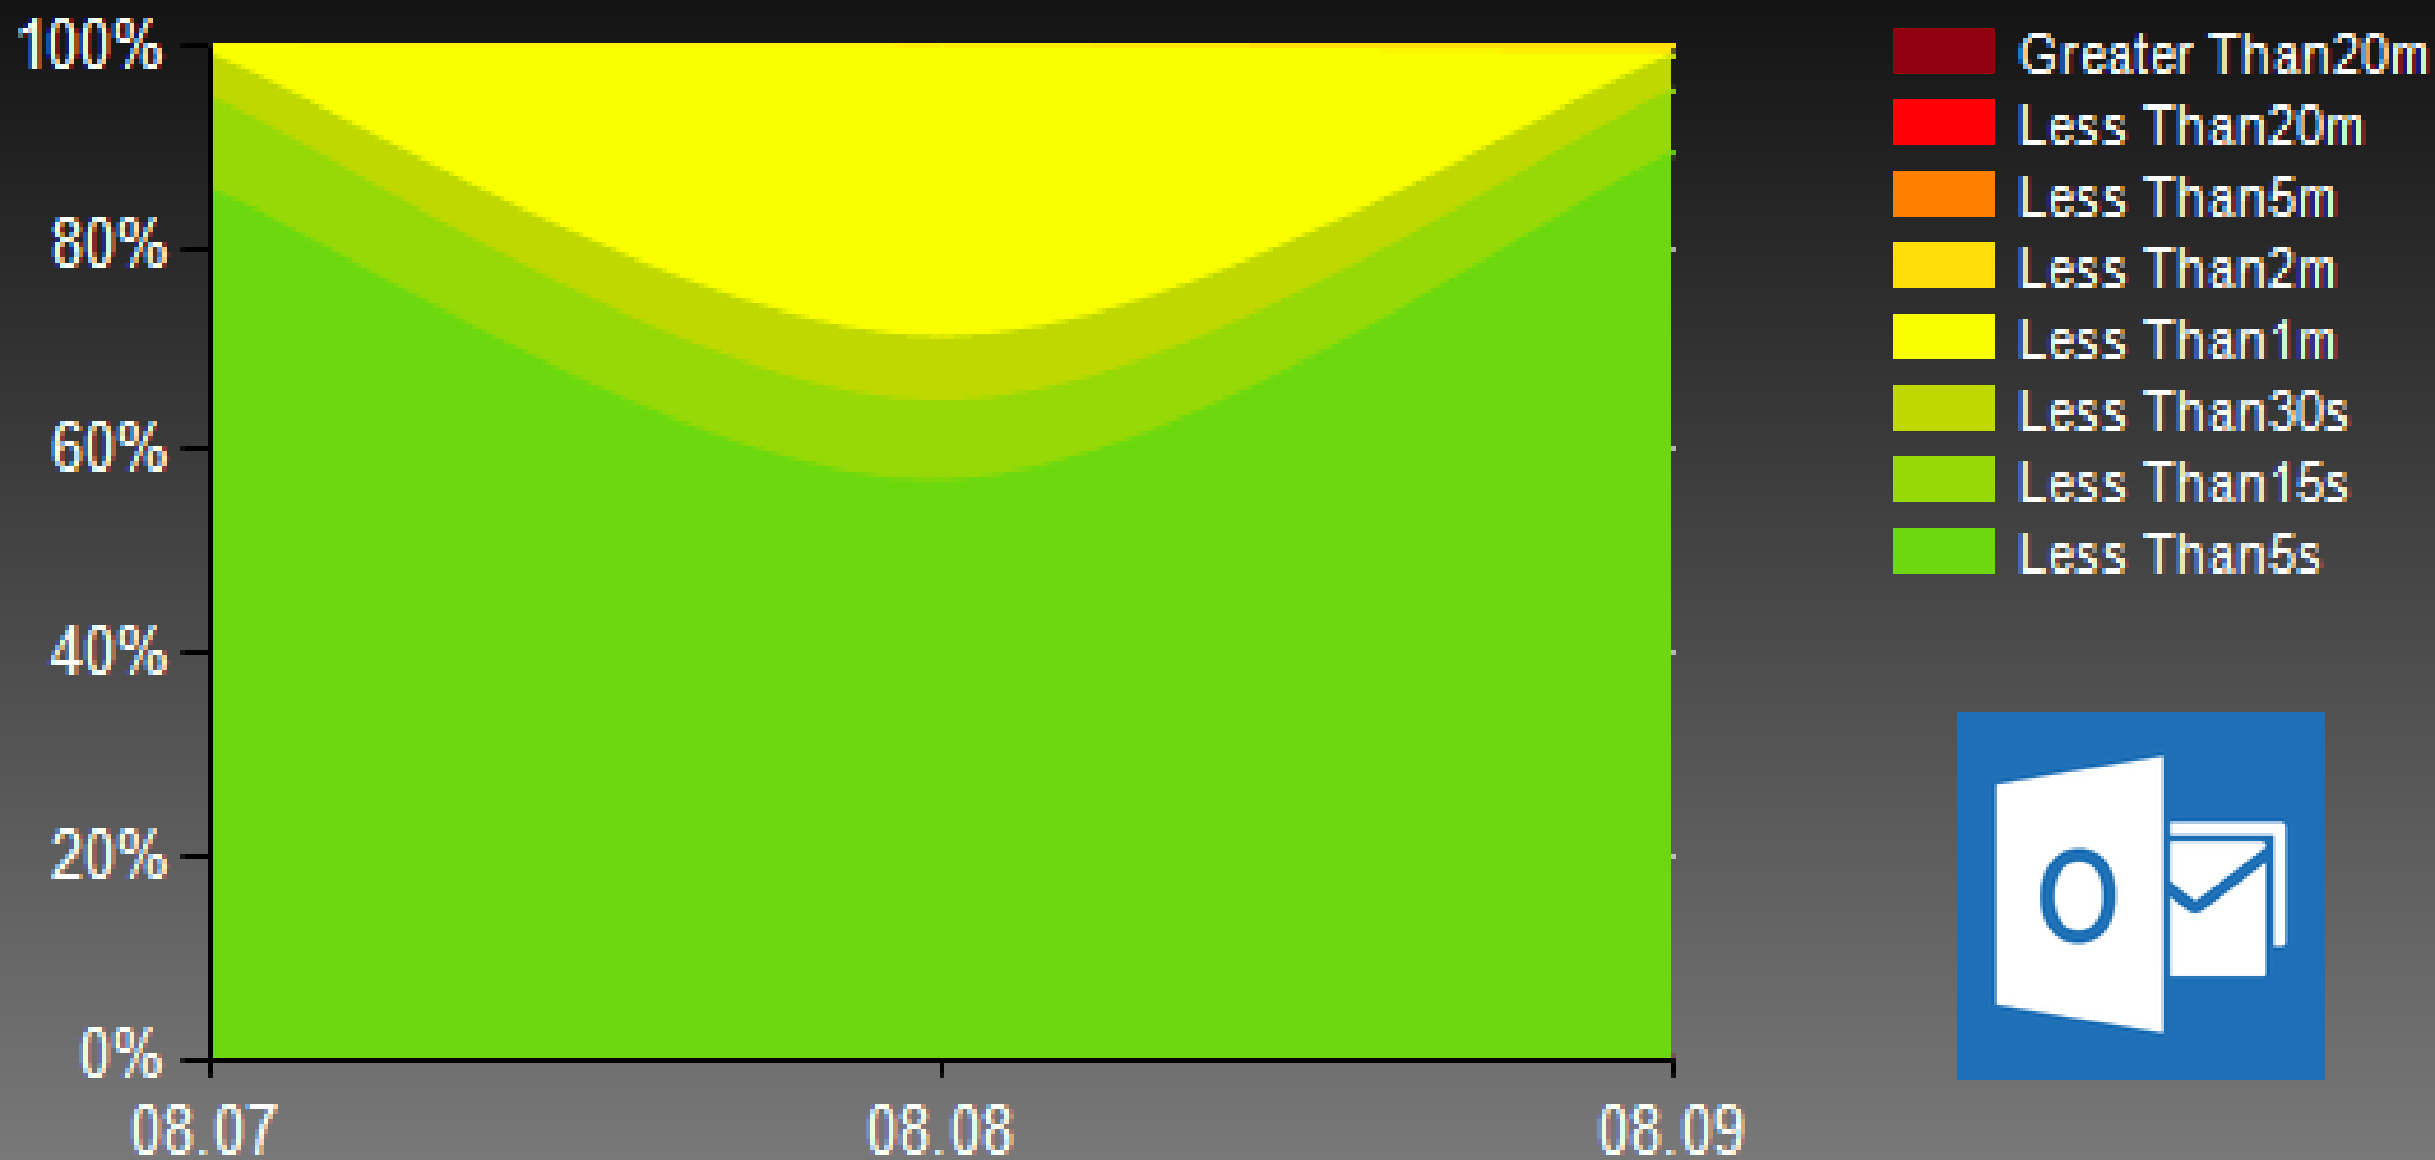
Make it easy

## Staged Data Acquisition






# Initial Connect: HTTP



# Your turn - Design for production-quality data



What might be your stages for risk mitigated data acquisition?



Role of active and passive monitoring?

How can you engineer for EaaS deployment?



Select your data sources

Determine the data necessary to  
answer your questions

## Information

 Online

Total Cluster : 6

Total Host : 45

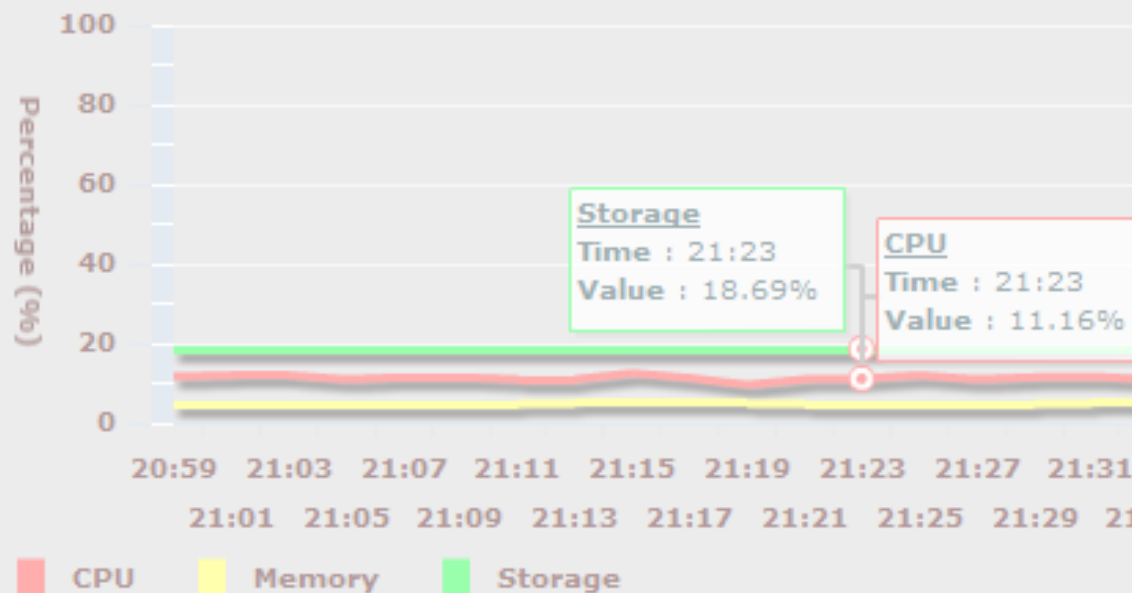
Total CPU : 502

Total Memory : 1.78 TB

Total Storage : 6.07 TB

Processes : 37

## Live Utilization Chart



Infrastructure  
data

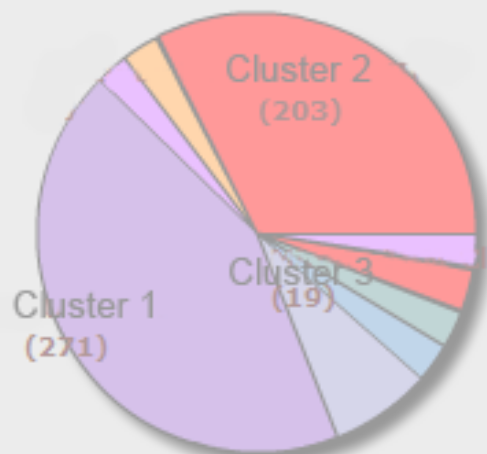
CPU

Memory

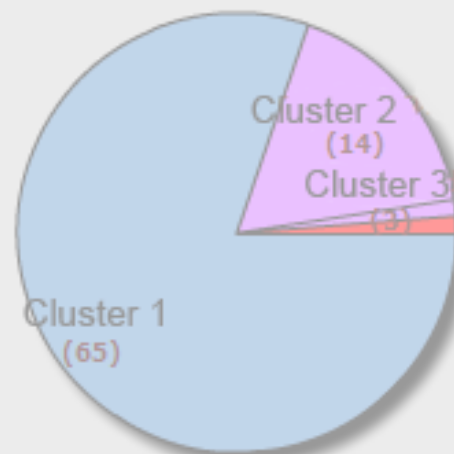
Storage

Network


## Event Statistic - Alert



## Host Status - Host Down



## Utilization - Storage

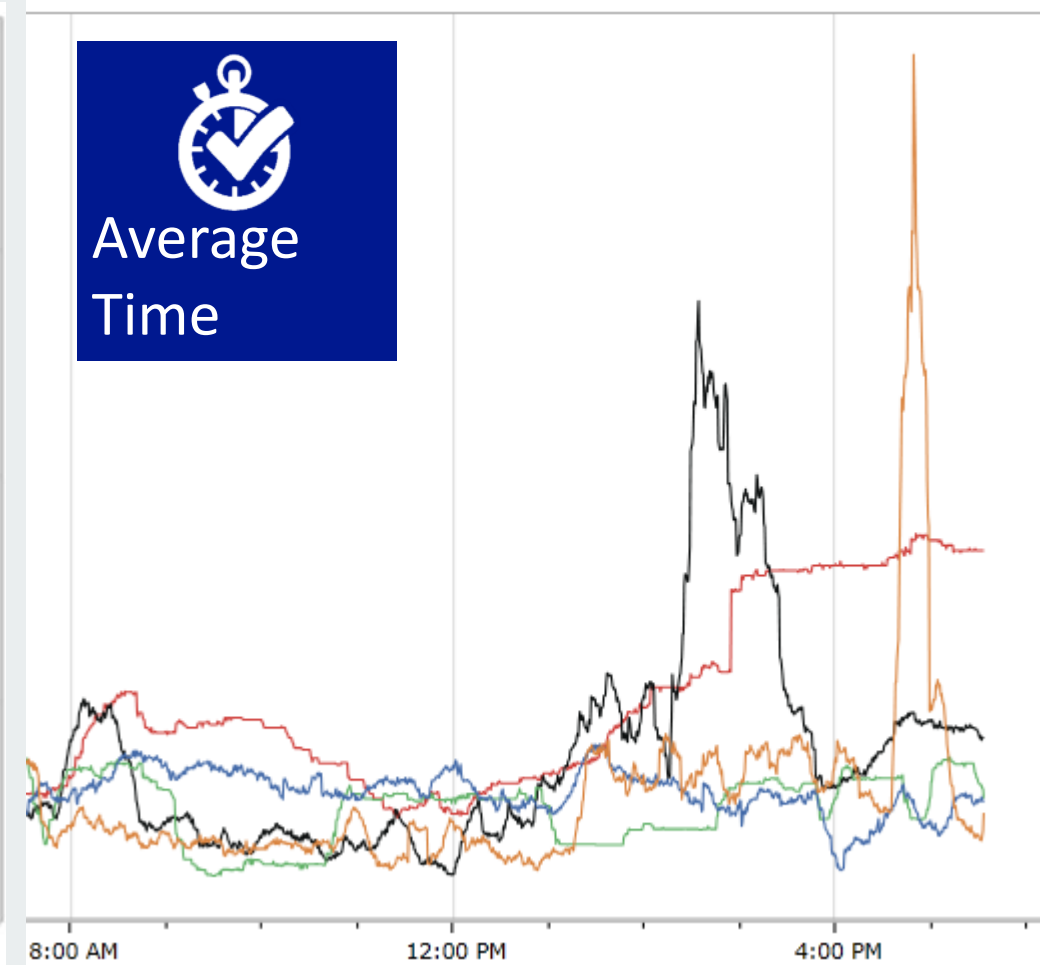
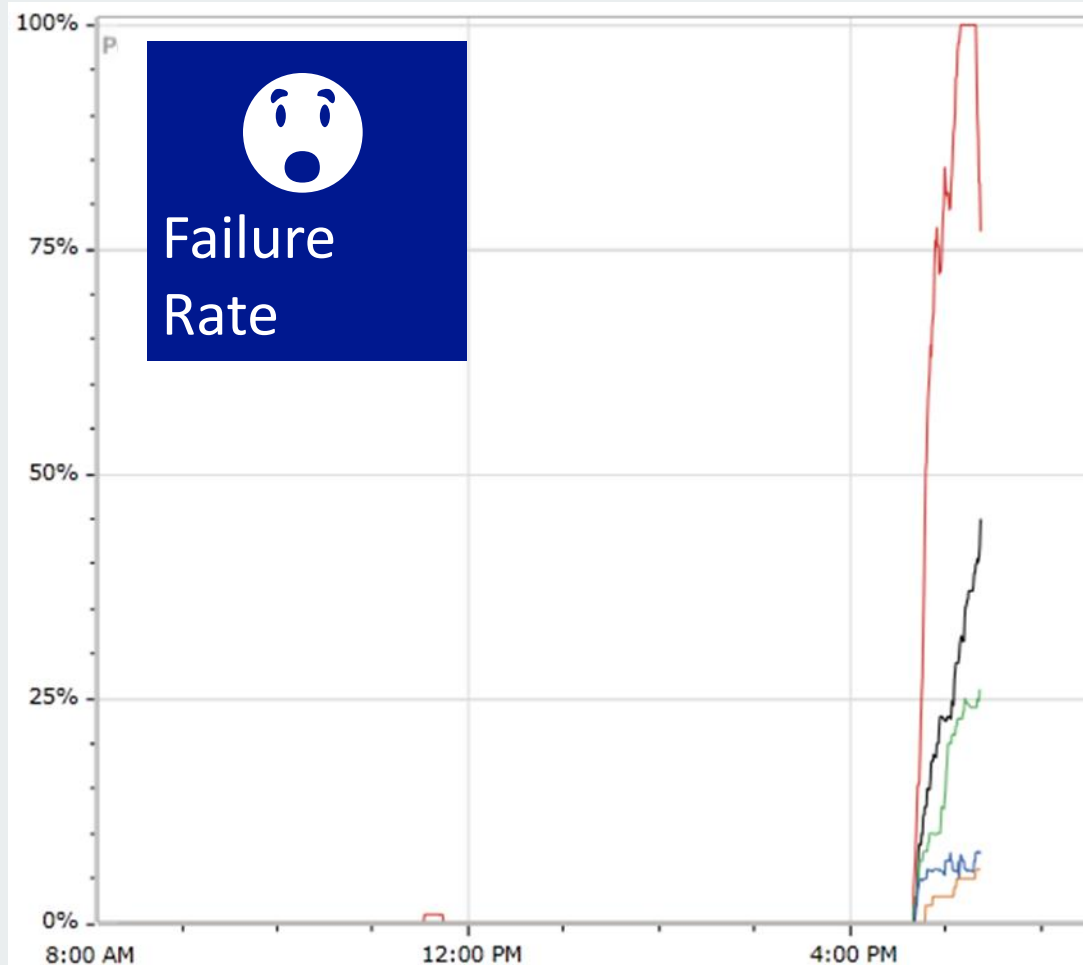


# Application data



New and  
Open

Download  
Template



# Hang and crash data

Specialized application data

Most frequently encountered conditions

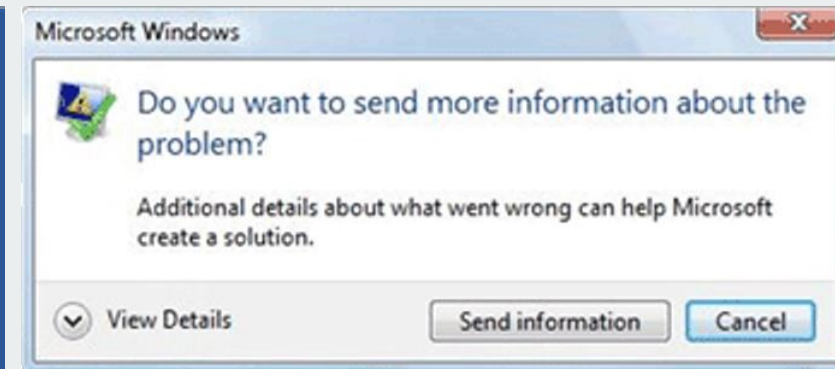
Bucket the Big Data and find the offenders

Get offending code and function calls (stack)

But it worked in Dogfood

Culprit was a old version add-in

Harden against this





# Usage data

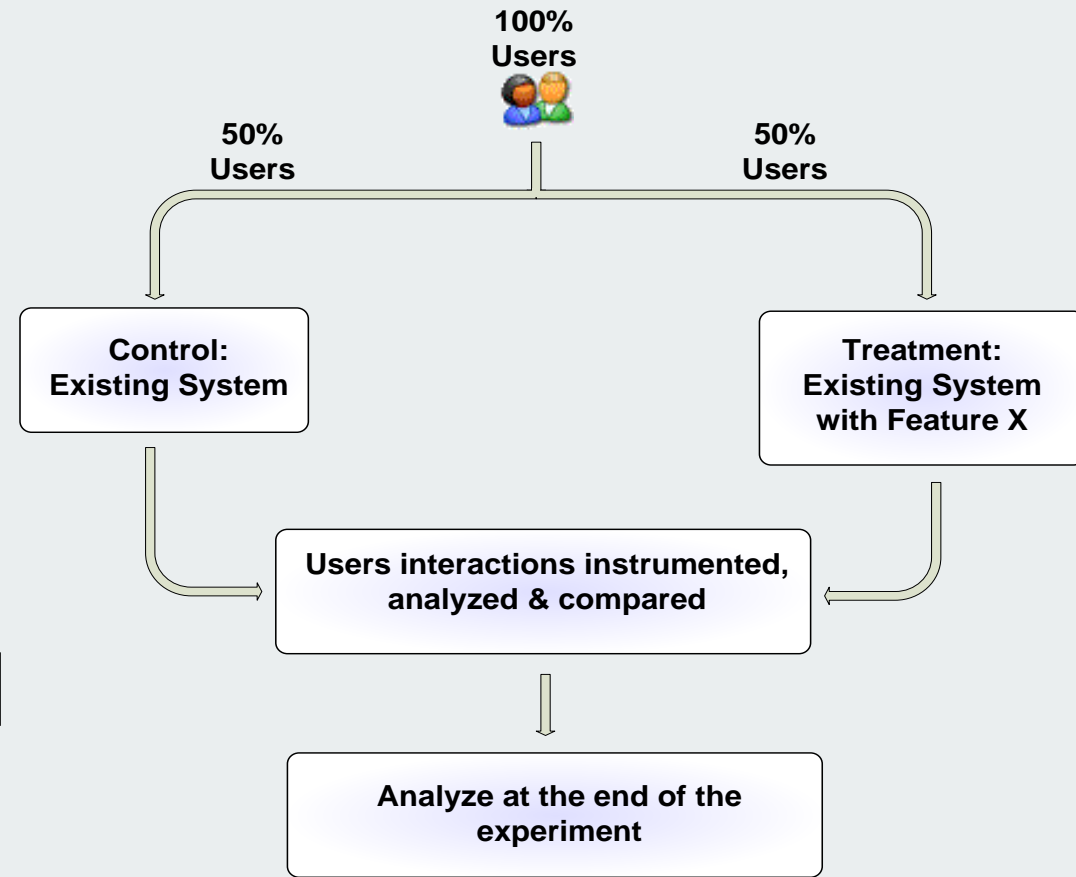


## Client-side instrumentation

Proprietary

Javascript: clicks, hovers [web apps]

Get 1x1 GIF: Page Views [web apps]



## Combine into more complex scenarios

How did user get to shopping cart checkout?





# Feedback data



Massive crack in dam



Oscars red carpet



I told you so



College queen busted



Russian troops



Oscar nominees



## Feedback about Bing

☒ Suggest ☐ Like ☐ Dislike

Enter your feedback here. To help protect your privacy, don't include personal info, like your name or email address.

400

Send



# Feedback data

**Bing Feedback** ✕

Click to highlight parts of the page you want to comment on. Hide areas you want to keep private.

**Highlight** Hide

Tell us more

Problem	Suggestion	Kudos
<i>Describe the problem, offer a suggestion, or let us know what we're doing right.</i>		

☒ Send email to the Bing team (sfeed)  
and cc me at  **Submit**

Search

Privacy Web Images Videos News Maps Shopping WolframAlpha Bing

Web Images Videos Shopping News Maps More | MSN Hotmail

**bing**<sup>™</sup>

Web **Microsoft** Finance Shopping News Wikipedia

RELATED SEARCHES

Microsoft Free Download Center  
Microsoft Office Download  
Microsoft Windows XP Updates  
Microsoft Internet Explorer  
Microsoft International  
Microsoft Templates  
Microsoft Program  
Microsoft Vista

ALL RESULTS 1-11 of 276,000,000

**Microsoft Corporation** **Great Result**  
[www.microsoft.com](http://www.microsoft.com) · Official site  
Get product information, support, and news from Microsoft.  
Windows 7 Internet Explorer  
Microsoft Update Download Center  
Microsoft Support Knowledge Base

Quick Access

Customer service 800-642-7676  
Search within microsoft.com

Financial 24.50  
▲ +0.10 (0.40%)  
US:MSFT

Product Wind  
Intern  
Office  
Xbox

SHARE Facebook Twitter Messenger Email

SEARCH HISTORY

Search more to see your history

See all  
Clear all · Turn off

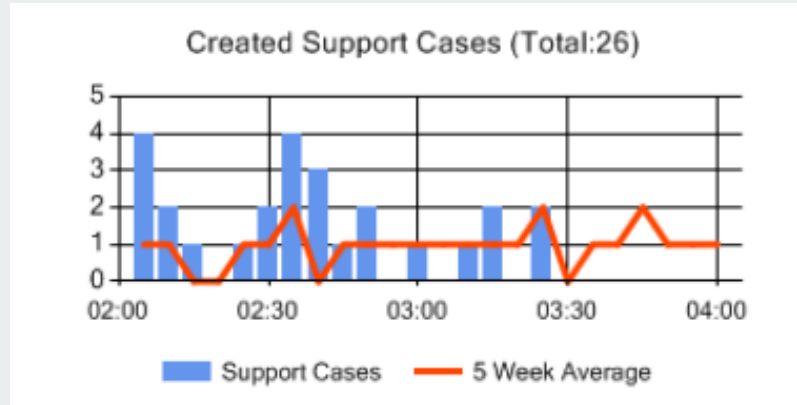
**Microsoft - Wikipedia, the free encyclopedia**  
History · Product divisions · Culture · Corporate affairs  
**Microsoft Corporation** (NASDAQ: MSFT) is a public multinational corpo  
headquartered in Redmond, Washington, USA that develops, manufactu  
supports a wide range ...  
[en.wikipedia.org/wiki/Microsoft](http://en.wikipedia.org/wiki/Microsoft) · Cached page

**Microsoft Support**  
The **Microsoft** Support home page is your support portal for **Microsoft** p  
Download updates and find top issues, error messages, and troubleshoo  
Windows XP ...  
[support.microsoft.com](http://support.microsoft.com) · Cached page

**Microsoft Europe** ...

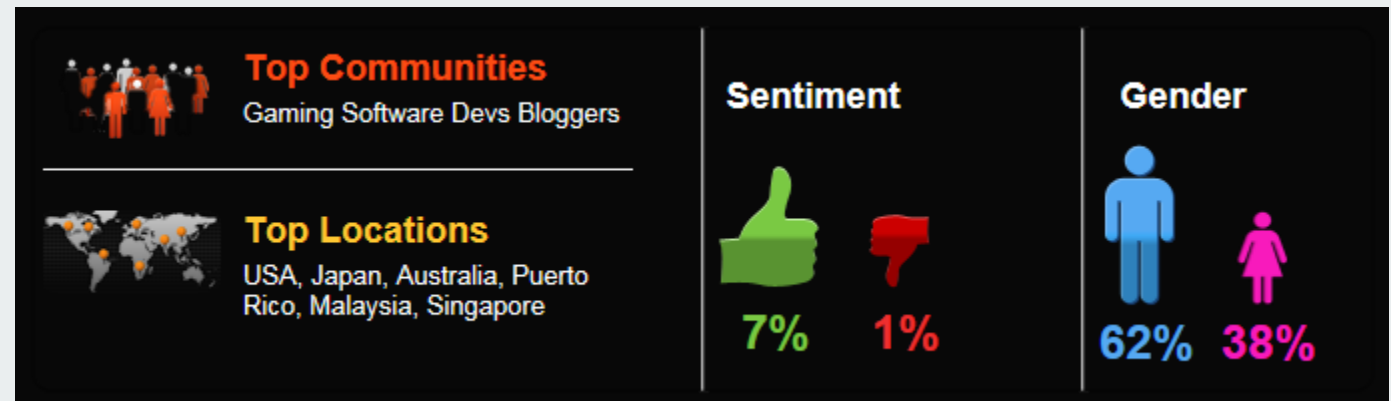
# Feedback data also includes...

## Customer Support Data

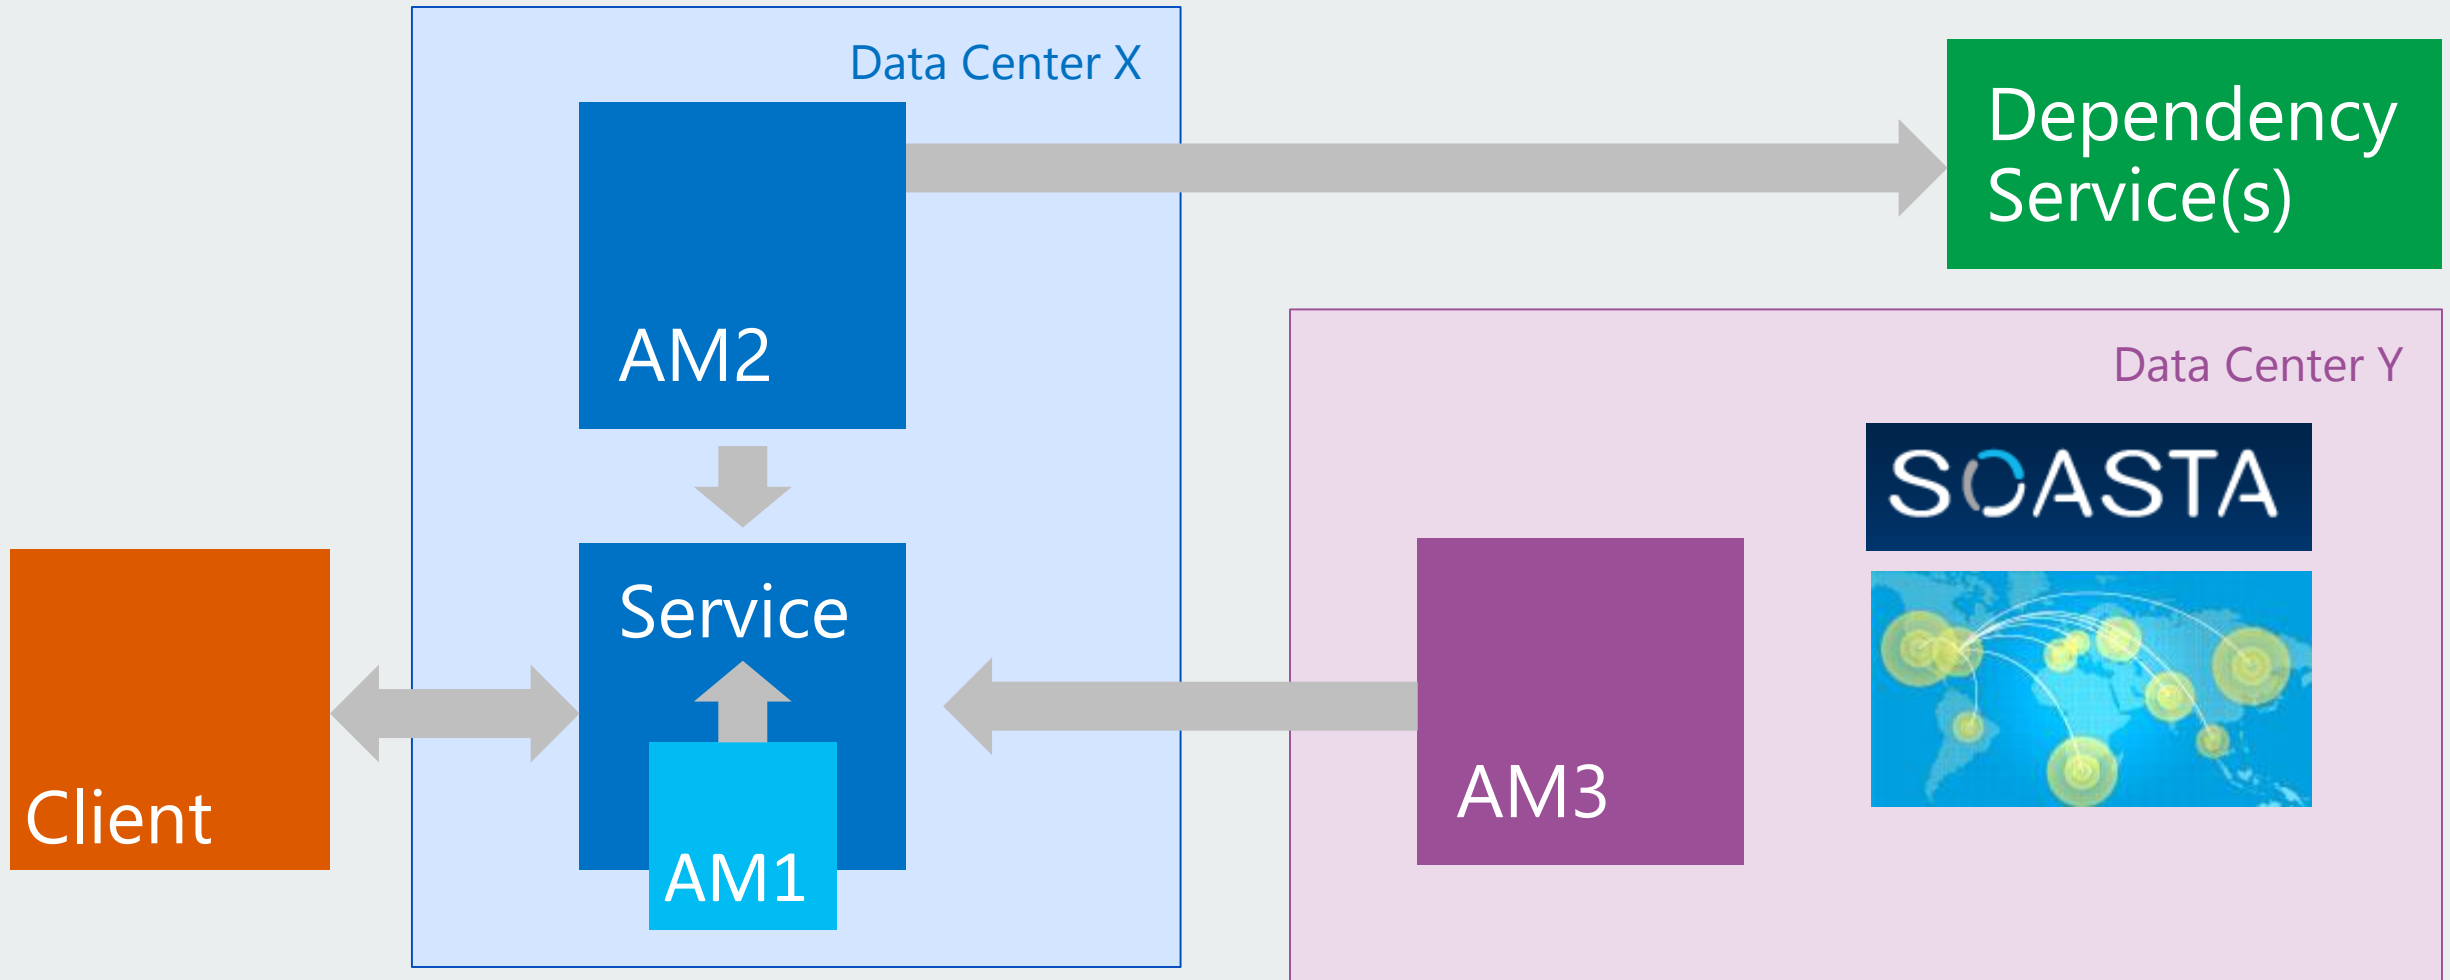


Service	Severity	Status
Lync Online	Sev1	➔ Service Degradation
Exchange Online	Sev2	🔄 Investigating


## Social Media Mining



# Active monitoring data



# Ease of detection



Again services have it easy

Many clients are always connected

On prem servers (enterprise) require partnerships

EaaS – Rich near real-time telemetry

# Data handling - privacy

Transparency and Control

Collection and Retention

## **Depends on Type**

anonymous Data

pseudonymous

personally identifiable info (PII)

sensitive PII

## **Depends on Purpose**

provide the service

improve the current service

improve a future version service

improve non-associated services

content personalization

ad targeting

# Data handling – non-service products

## Client and on-prem server considerations

### User owned resources:

bandwidth

battery

disk, cpu, etc...

### Correlations

end-to-end across clients and services

by user , by session



# Your turn - Select your data sources



Infrastructure data

Application data

Hang and crash data

Usage data

Feedback data

Active monitoring data

Design to handle your data





# Use the right data tools

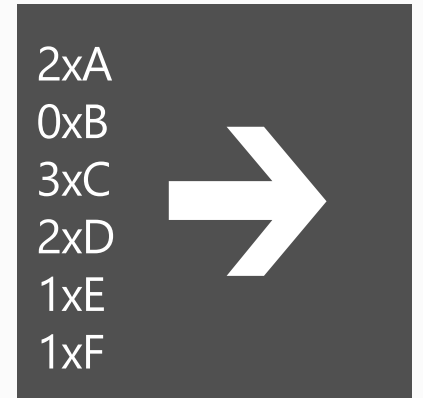





## A 50,000 foot view



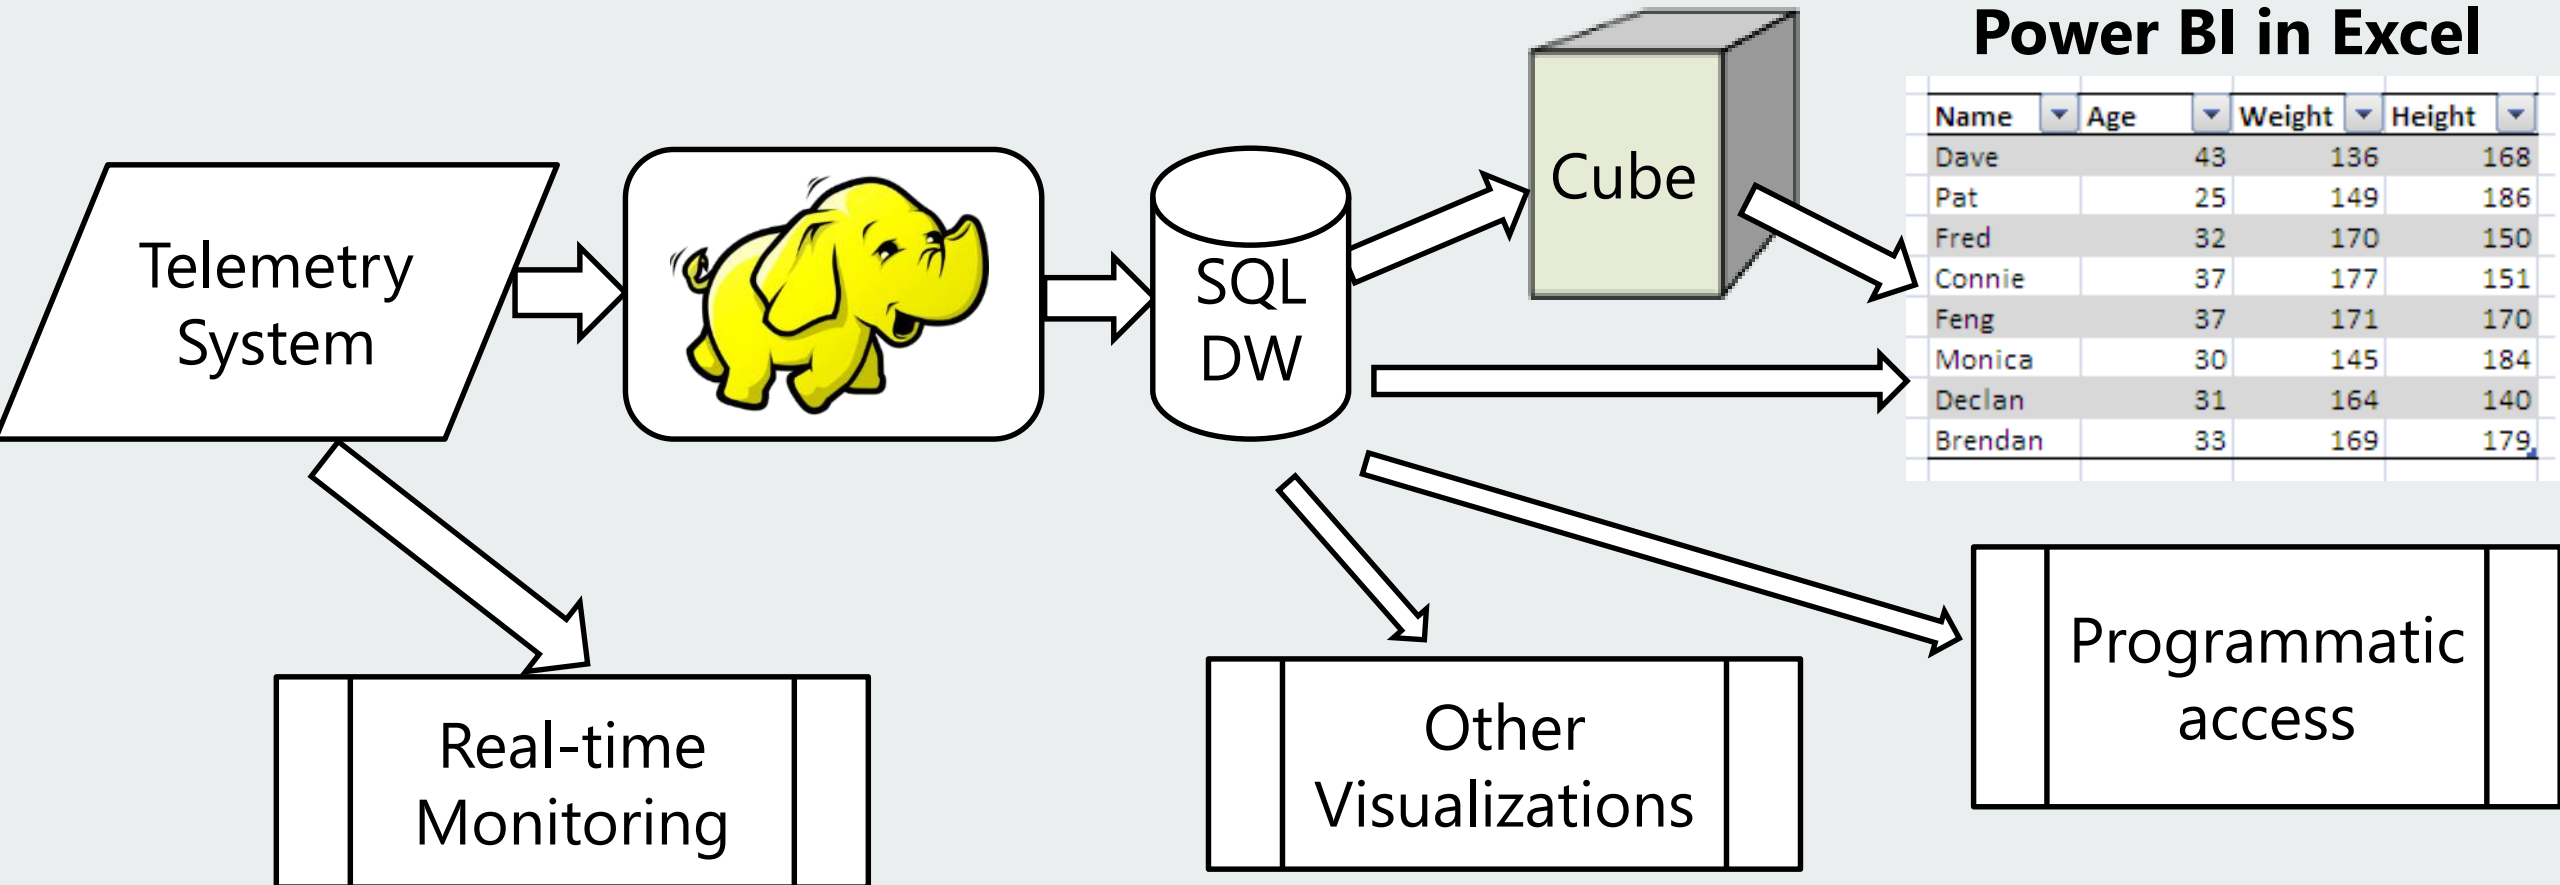
# Data storage and processing systems


Database	Table storage (SQL)	Optimized for CRUD – (create, read, update, delete) of single records
Data Warehouse	Table storage (SQL)	Table structure optimized for queries and bulk insert
OLAP Cubes (Online Analytical Processing)	Multidimensional Cube (MDX)	Aggregations of measures for multiple dimensions
Hadoop / Map-Reduce	Distributed File System (HDFS)	Big Data TB-PB-EB Unstructured Data

# Hadoop in 60 Seconds














# A common data flow





Pinned	Color	Metric	Instance	Aggregation	Value	DisplayedSum	DisplayedAvg	T
<input type="checkbox"/>	<div></div>	SessionsByDatacenter.FailureRate	CO1		0.07		0.08	2

# Feedback



# Your turn - Use the right data processing tools




DB, DW, Cube, Big-Data platforms

Put it all together



Do you need real-time monitoring?



...sentiment analysis?






Get answers to your questions  
and learn new questions

# Outlook.com prioritizes performance

View Inbox – Page Load Time (PLT) by Browser  
As experienced by actual users




JSI  
JavaScript  
Instrumentation

500 Million  
measurements  
per month



Availability




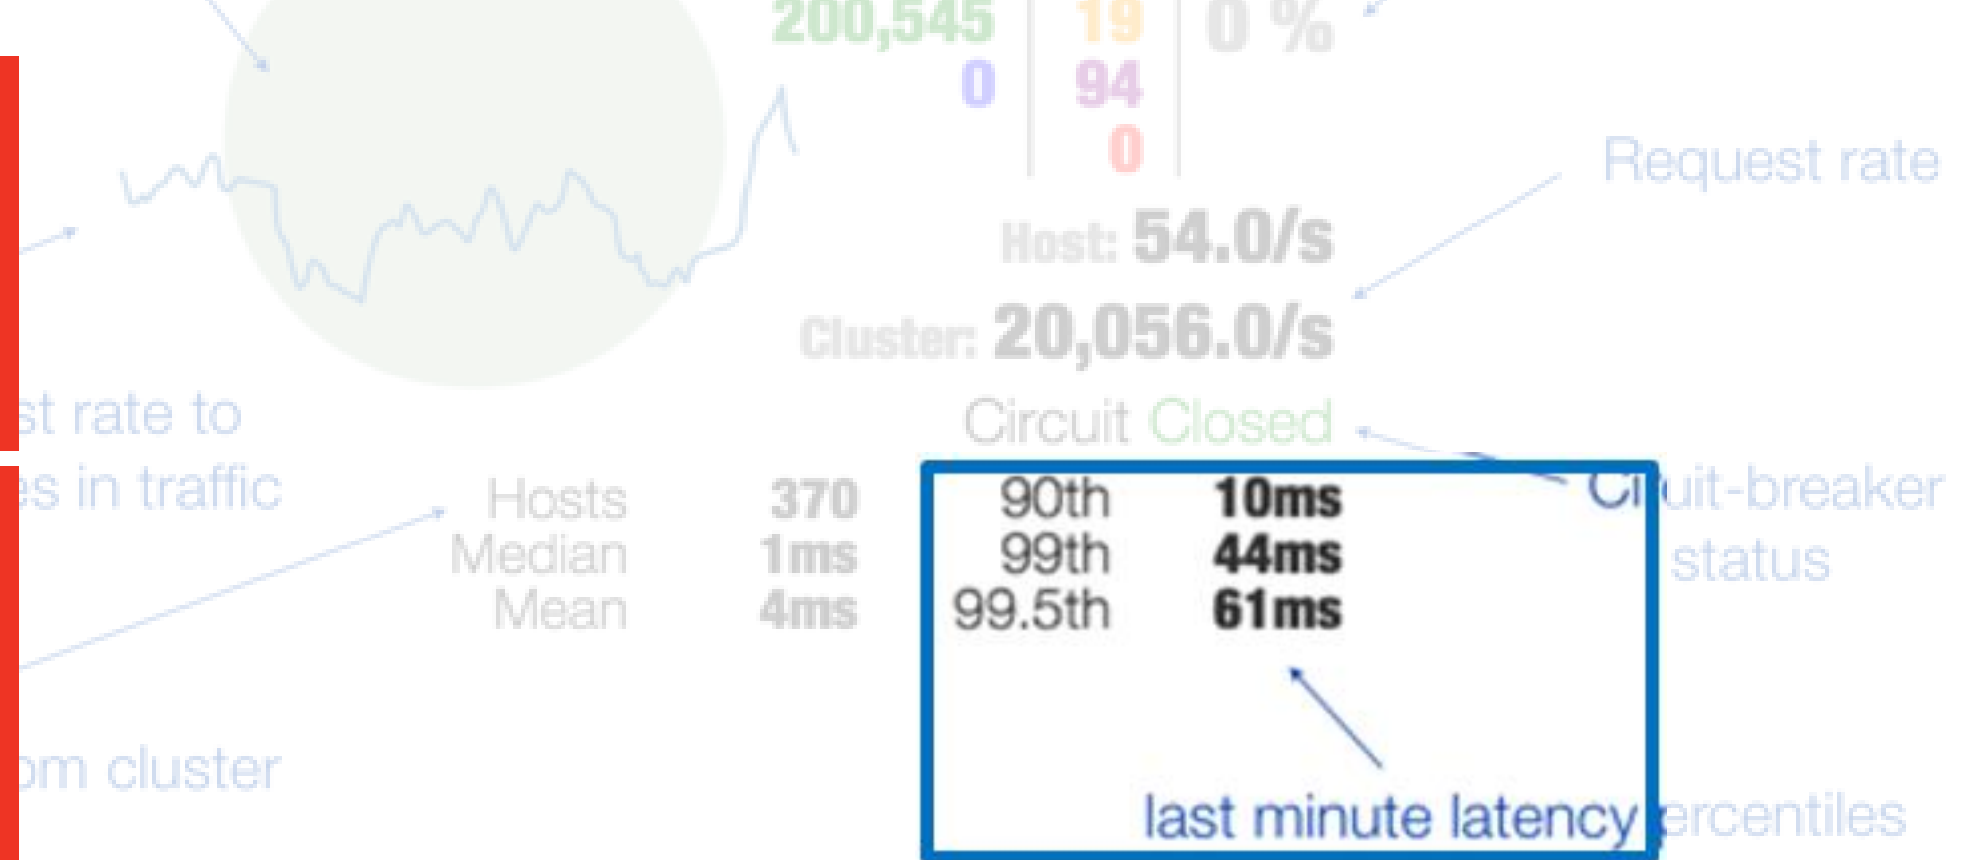
Prioritizes  
availability

Predict 75% of  
dips 24 hours  
ahead of time

Time

# NETFLIX

Netflix prioritizes  
perceived  
availability &  
performance



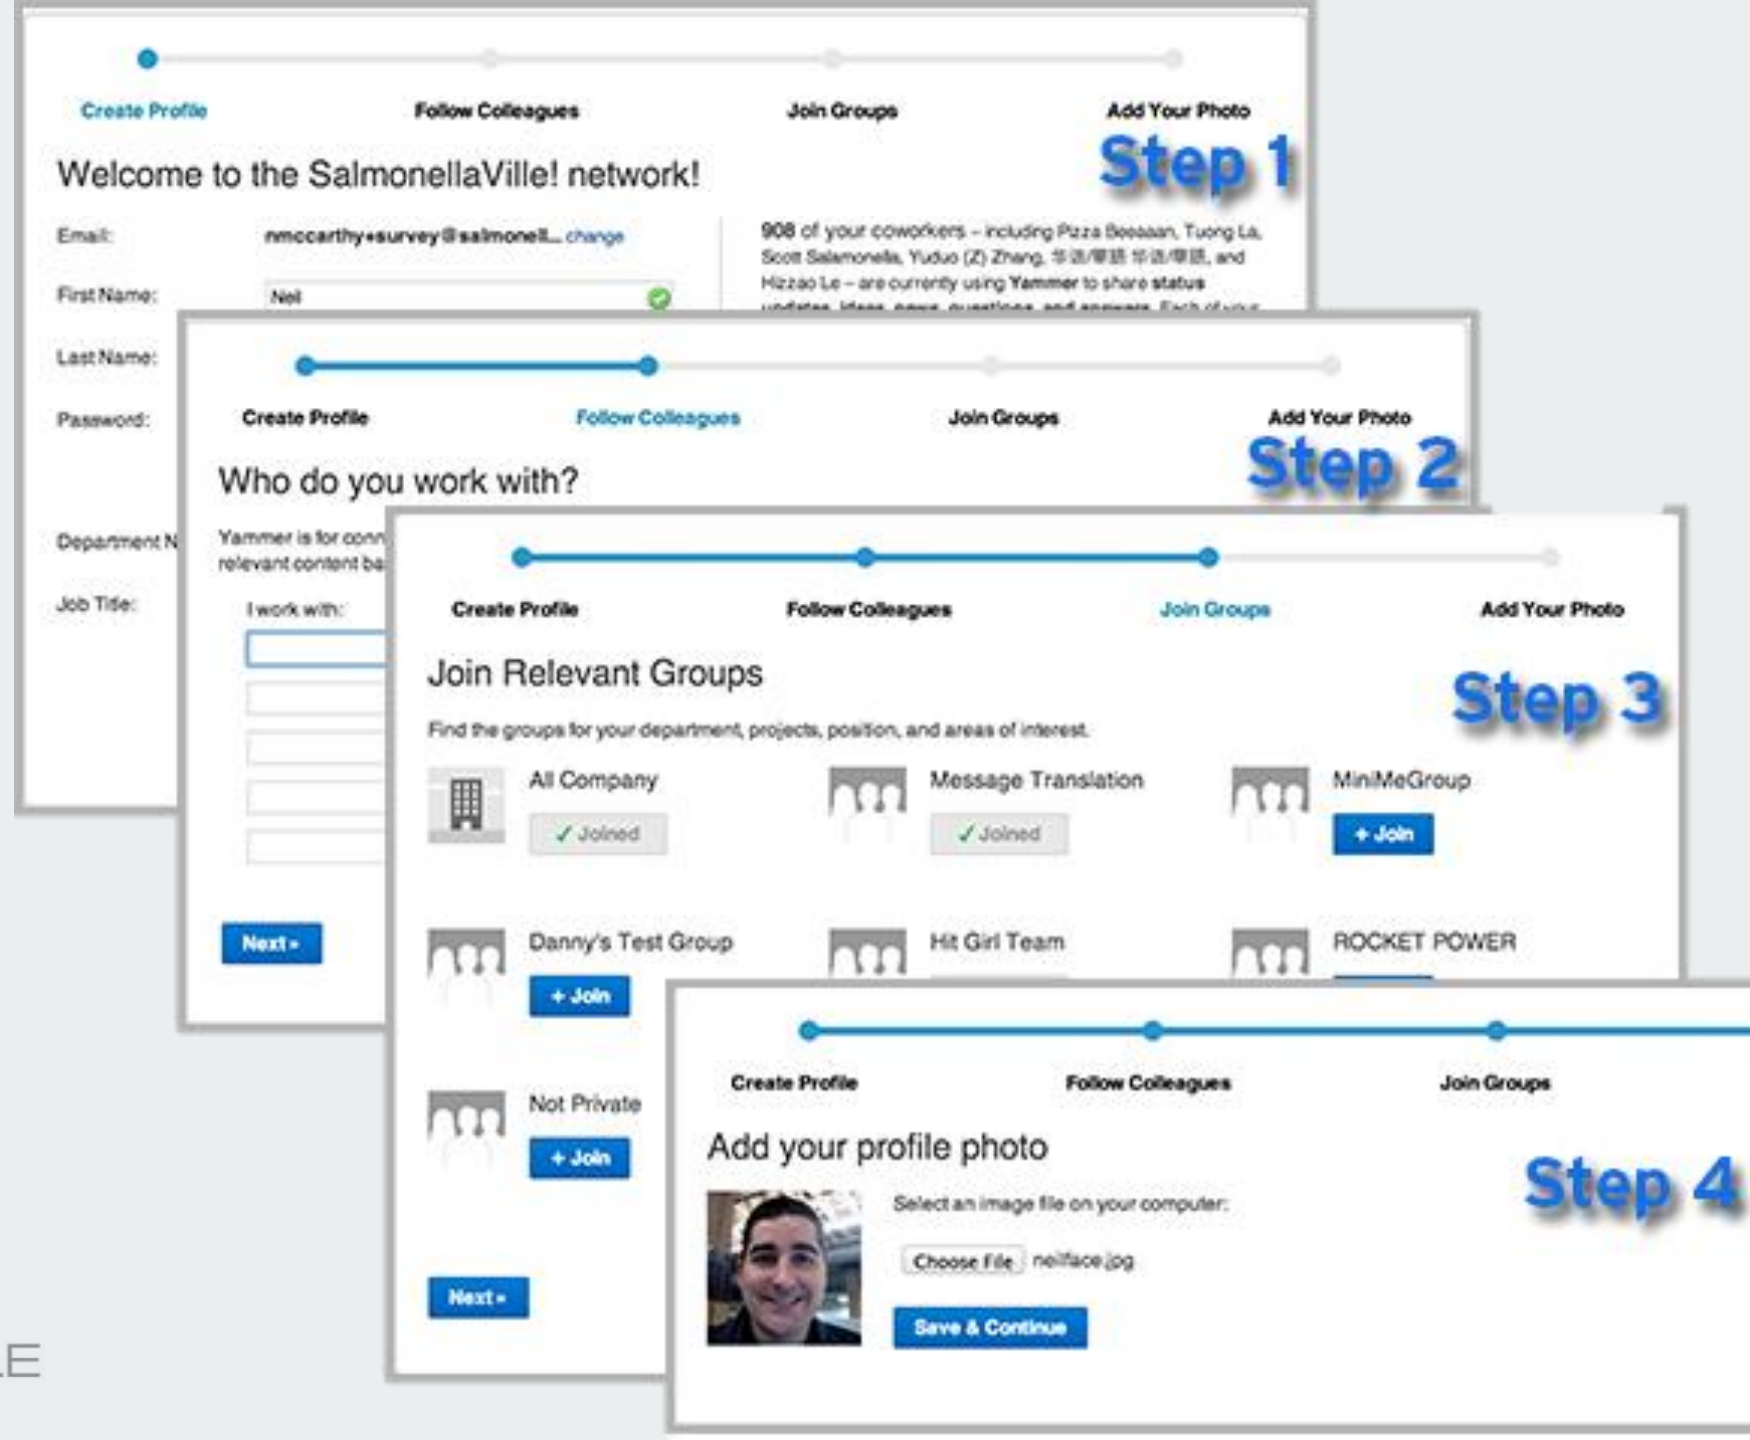


Yammer  
prioritizes  
usage

What happens to new  
user retention when  
you shorten the signup  
flow?

It goes down!

Don't ship that feature




# Find new questions and repeat



Try algo  
Collect data  
Adjust algo  
Collect data  
repeat



# Find new questions and repeat




In Visual Studio 2012  
we added asynchronous  
loading for *solutions*



We also  
added  
Telemetry

Results look  
good





# Your turn – Get Answers, more questions



Where are the answers to your quality assessment questions?





# Summary



Determine your questions

Design for production-quality data

Select your data sources

Use the right data tools

Get answers to your questions

Learn new questions

Repeat

The most recent version of this deck can be found at <http://setheliot.com>

# Special thanks to these folks

Ravi Vedula

Andrea Jesse

Bill Hodghead

Danny Thayer

Joseph Sefair

Kitty Thomas

Amanda Reinke

Heather Lader

David Brooks

Mike Tholfsen

Jodie Draper

Brian Mueller

Tara Roth

Dror Cohen

Nathan Halstead

Lori Oviatt

Monica Catunda

Lynette Skinner

Joe Schumacher

Donny Luu

John Hoegger

Alain Anyouzoa

**Any questions?**



**ALM**FORUMSEATTLE

April 1-3, 2014 • Washington State Convention Center

**Thank you for joining us!**