

www.STPCon.com

April 14 -17 2014



Create Your Roadmap to Data-Driven Quality


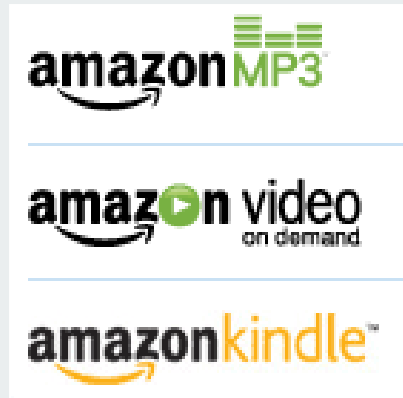
Seth Eliot

Principal Knowledge Engineer, Test Excellence



April 15, 2014

About Seth




Digital
Media
Services

A/B Testing
of Services

Petabytes
Processed

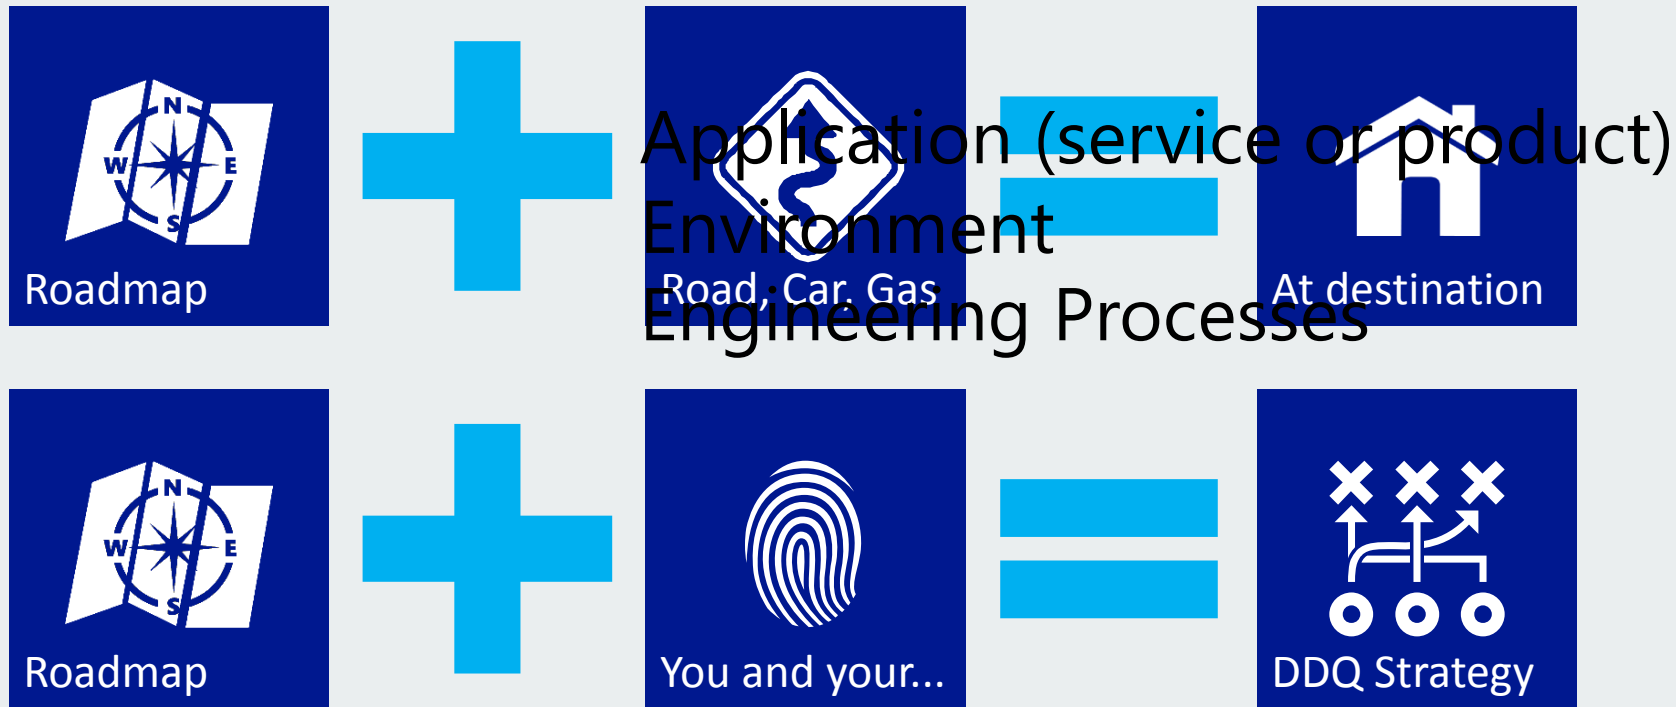
Services and
Cloud and
Data-Driven Engineering
and...





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



Highest Paid
Person's Opinion

Engineering Data

Test pass/fail results

Bug counts

Delivery cadence

Code coverage

Code Churn

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


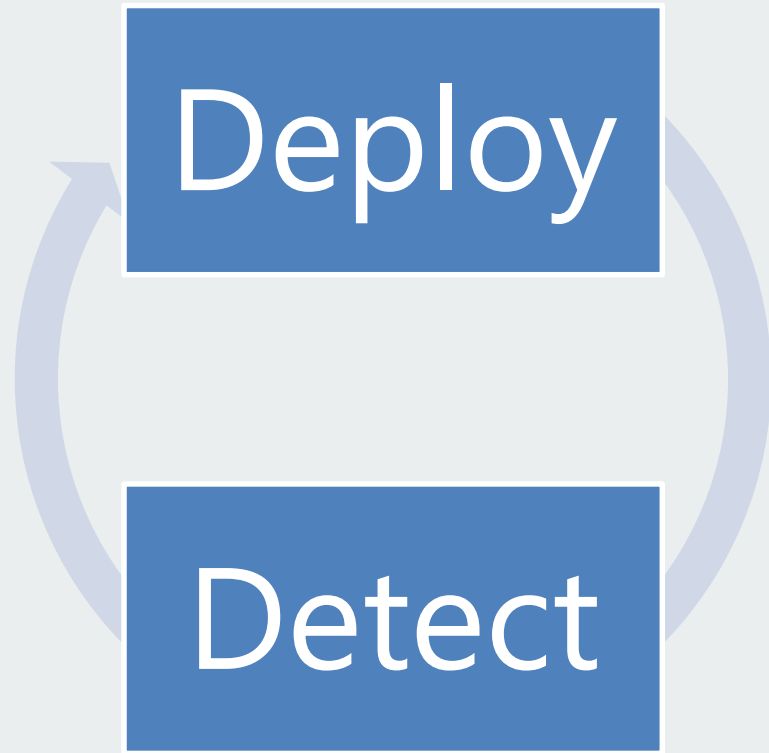




“If you have to kiss a lot of frogs to find a prince, find more frogs and kiss them faster and faster” - Mike Moran



Virtuous cycle of DDDQ



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



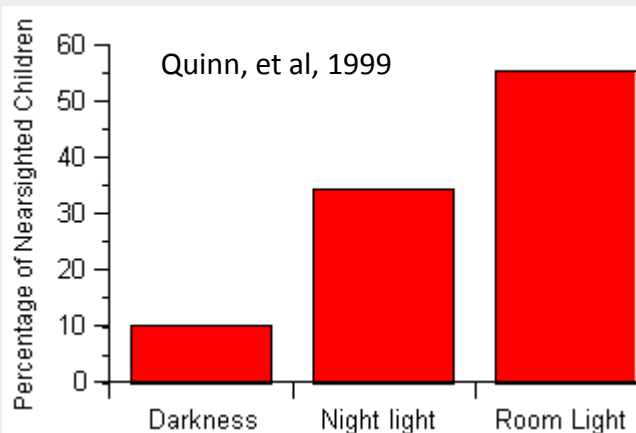
Determine your questions

What do you need to know about quality?

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”

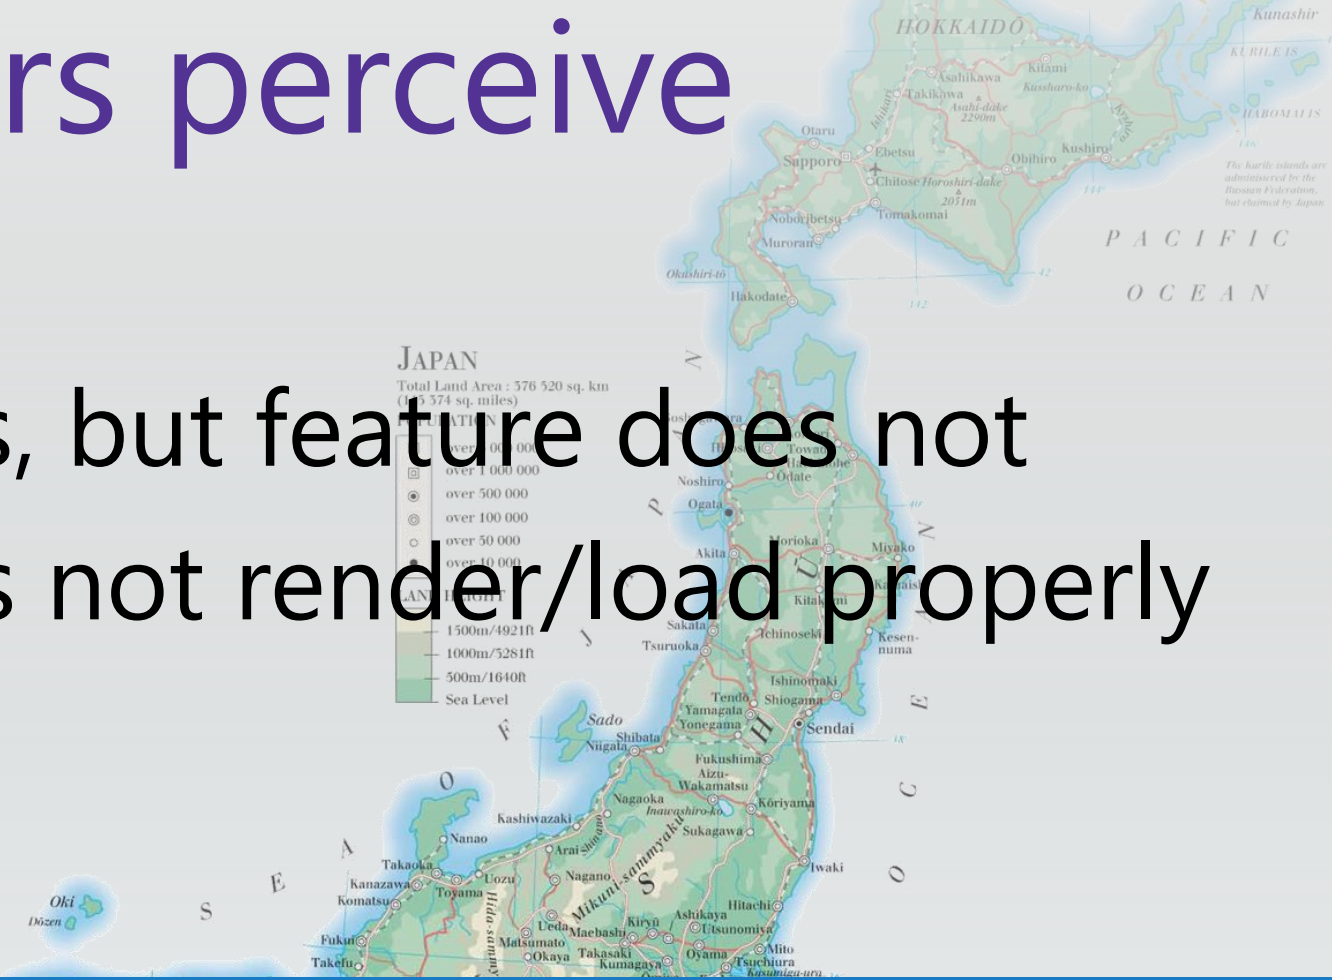


Critical work
stoppage for
user

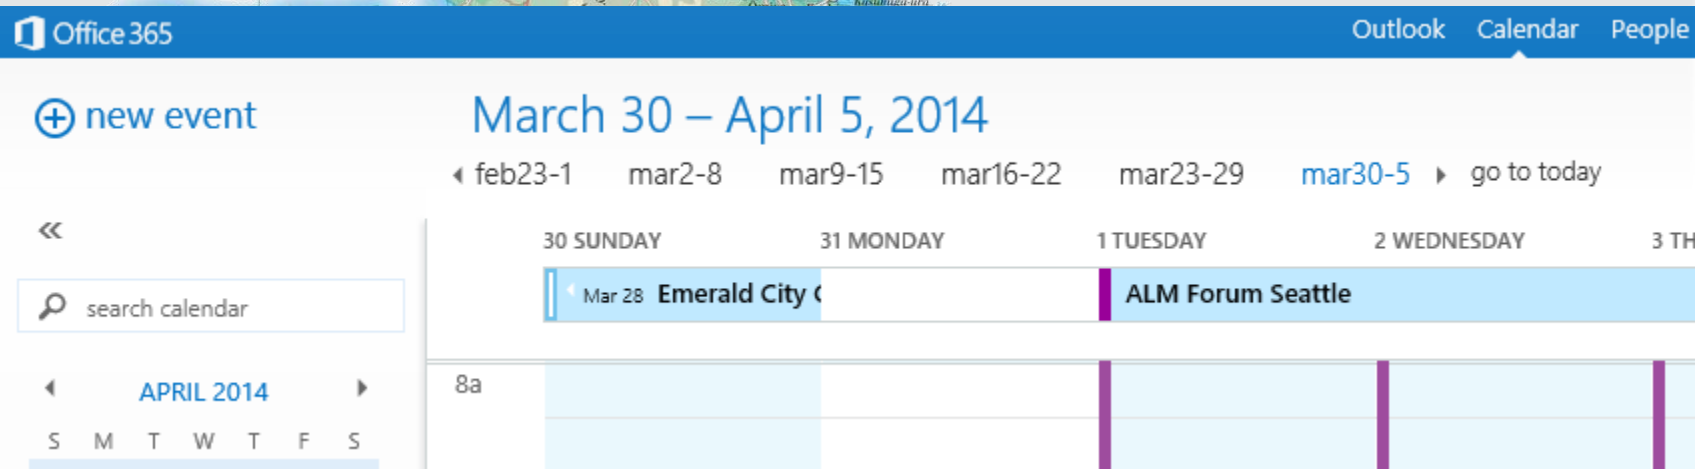


How do users perceive availability?

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

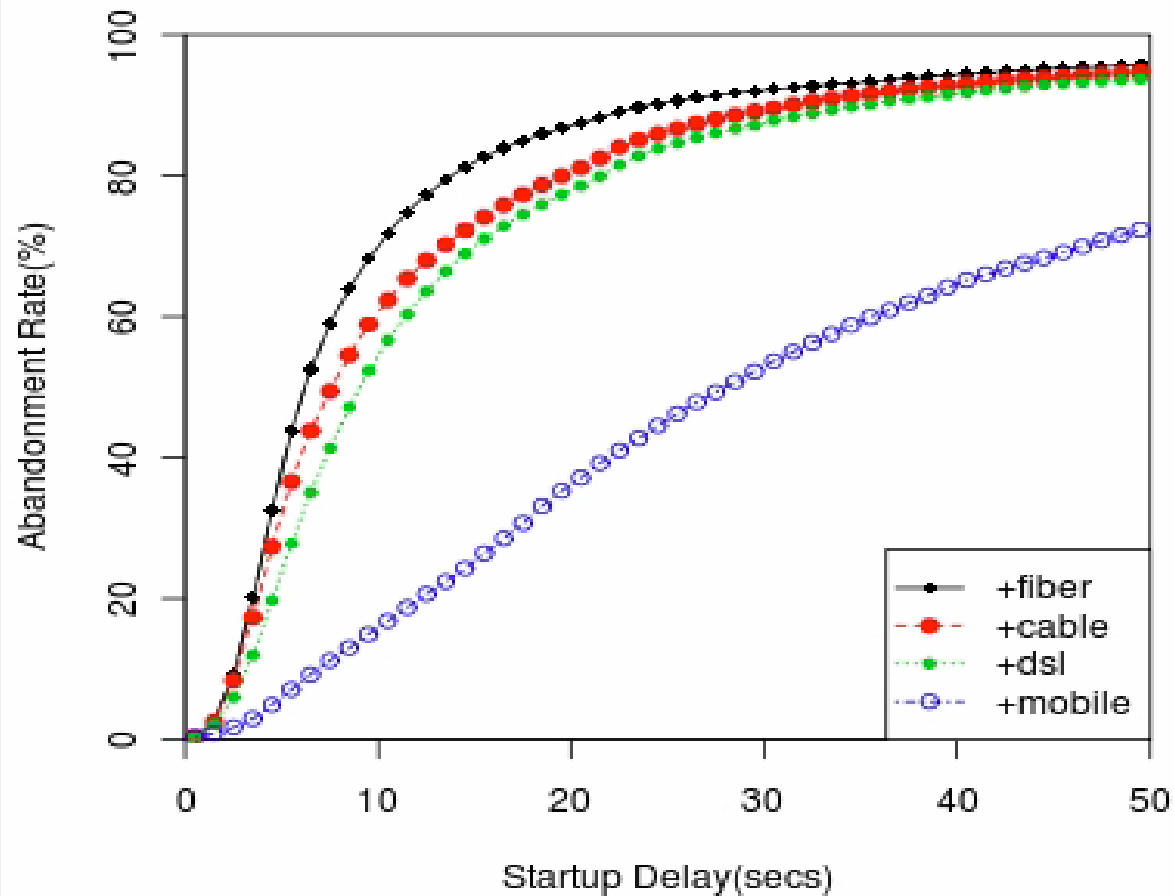


違った言語




How is performance?

Huge Impact



Power of Production Data



Real users



Multiple environments



End to end



Scale & geo-diversity

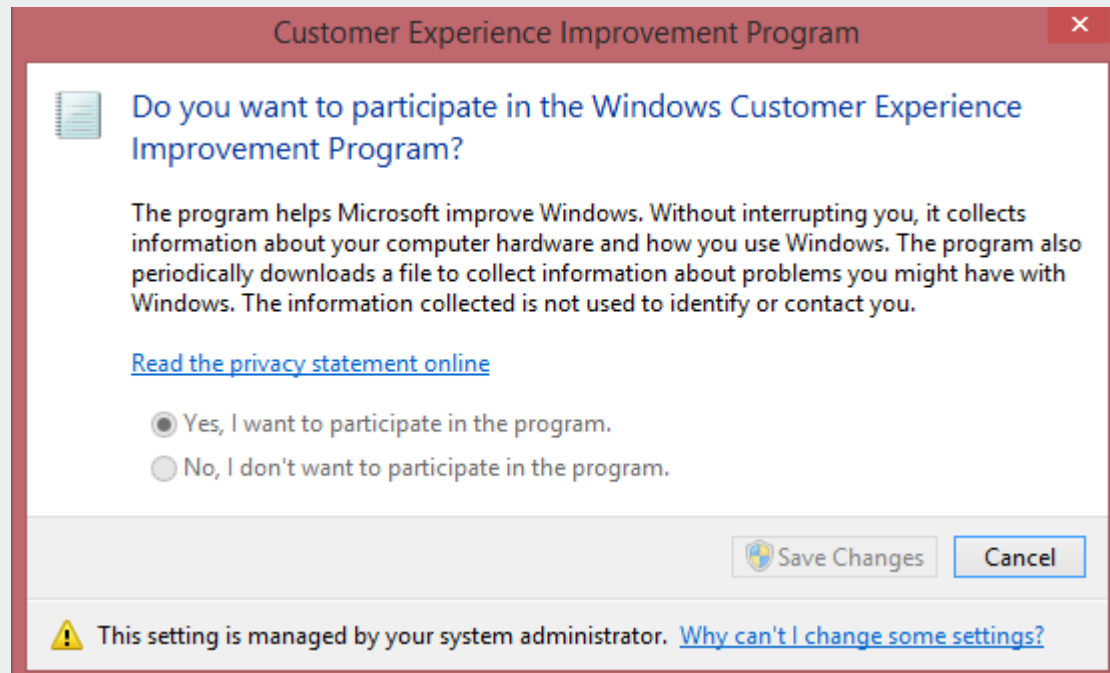
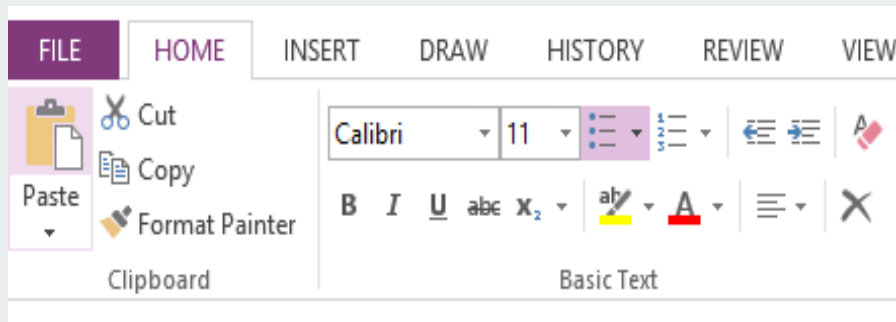
What do users do?



Customer Experience Improvement Program (CEIP)



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?

BROWSE FILES LIBRARY

EE Data-Driven Quality Validation EDIT LINKS

Documents > Office DDQ - 2013

Home Notebook Documents Recent

new upload sync edit manage share

All Documents Ready For Review (by EE SME) Find a file

✓	Name	Modified	Modified By	Checked Out To
	Completed Action Plan	February 02	Seth Eliot	



What is high priority?



Design for production-quality data

Get data from or near production

Two types of data to acquire

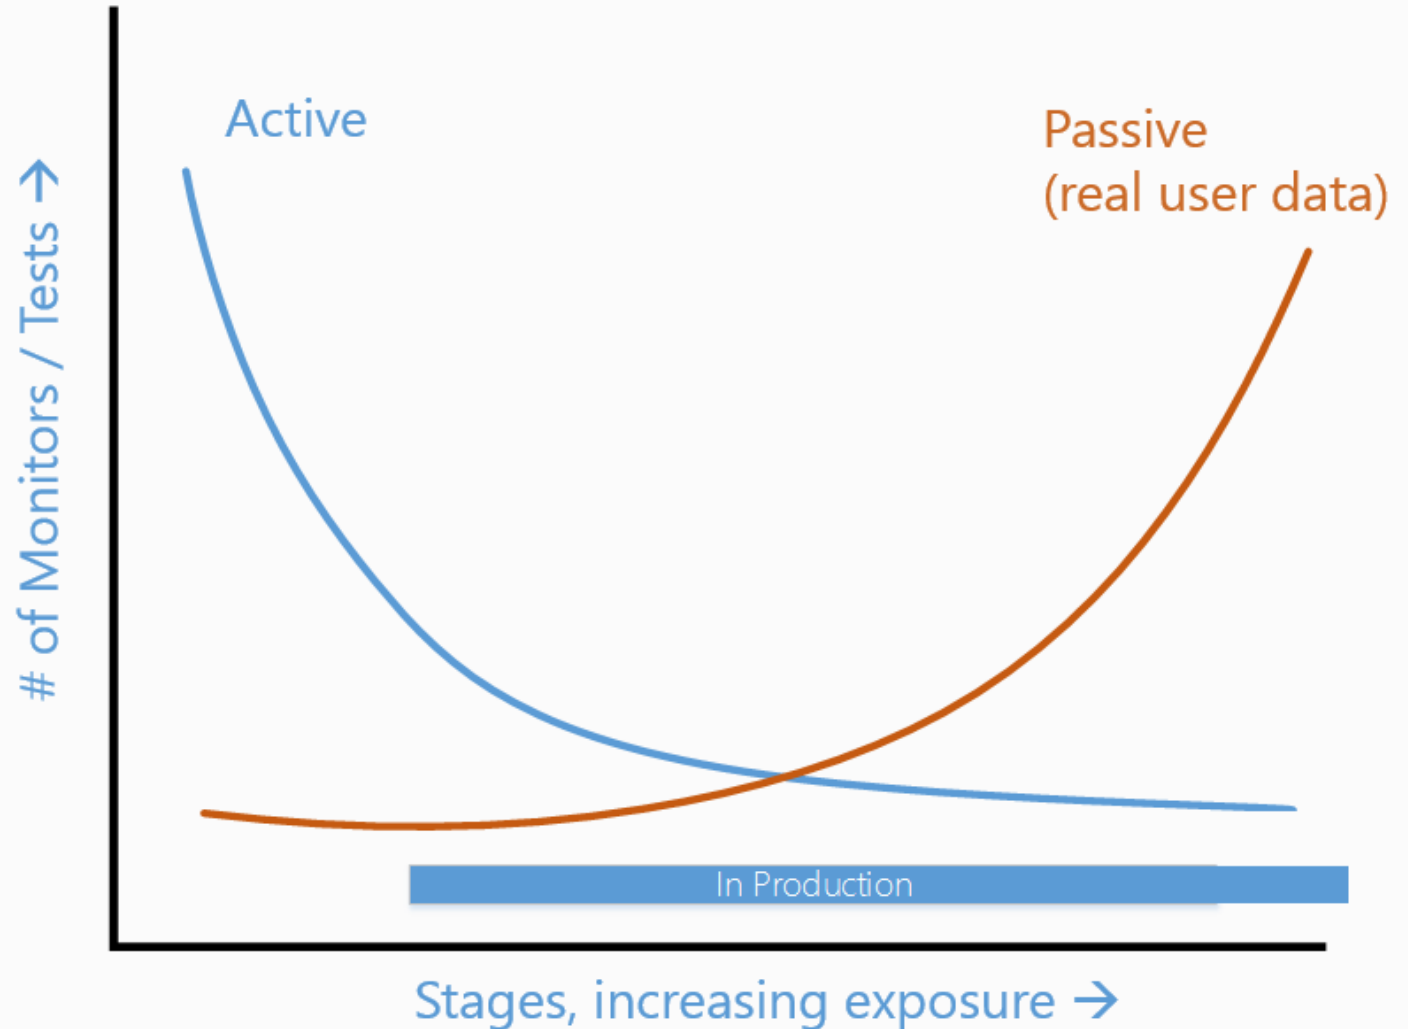
Active: synthetic

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)



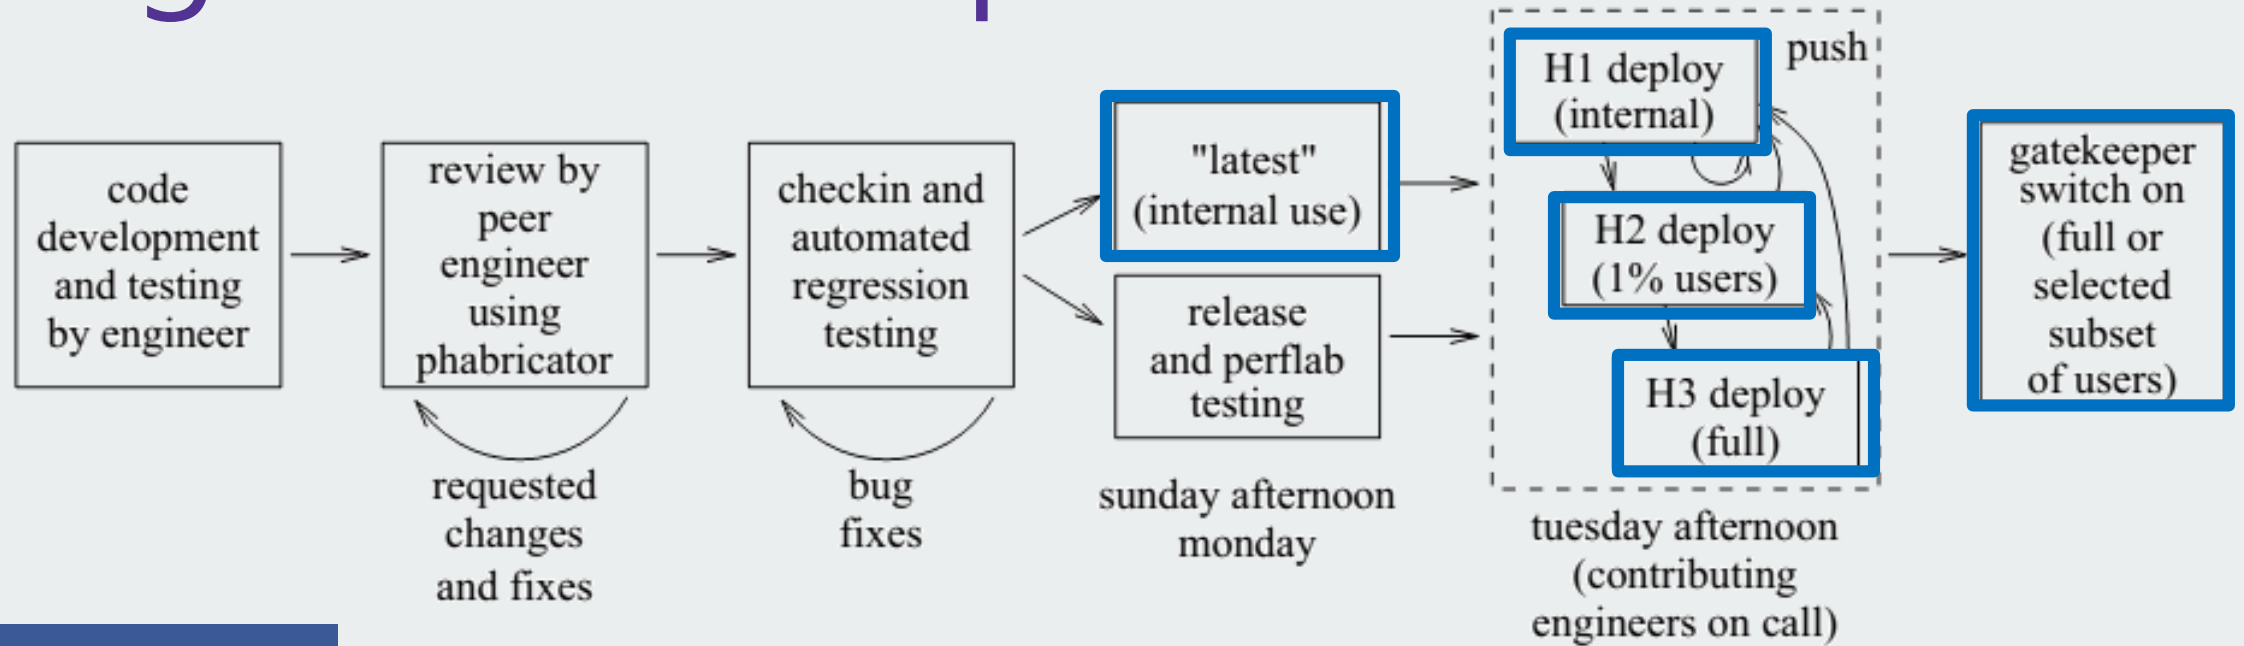
Staged data acquisition - Netflix



1B API requests per day



Staged Data Acquisition - Facebook



facebook

Dogfood

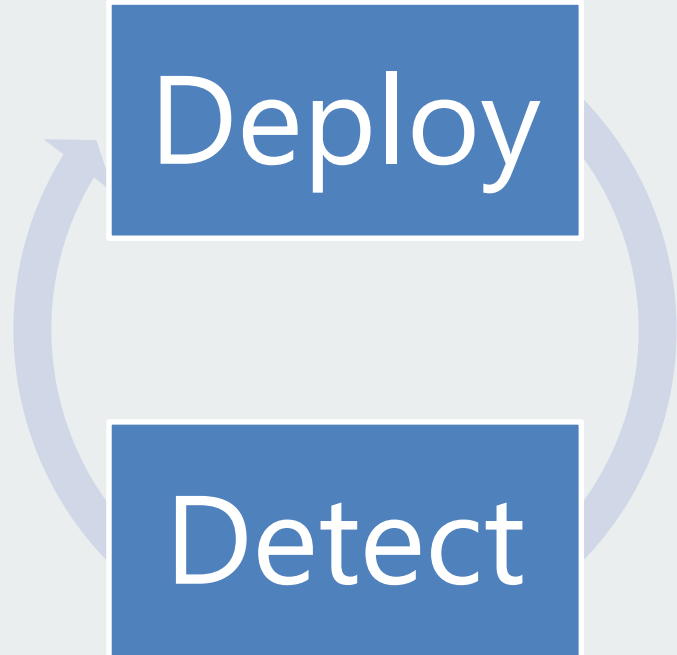
In prod, no users (except internal ones)

Some servers in Production

World-wide deployment

Feature flags

Speed of deployment



Usually easy for services

Client apps may have a deployment capability

Or may make use of feature flags

EaaS - Everything is now connected and thus updatable

Data acquisition for clients

Filtering and aggregation at client

Be kind to the client

Don't abuse user resources:

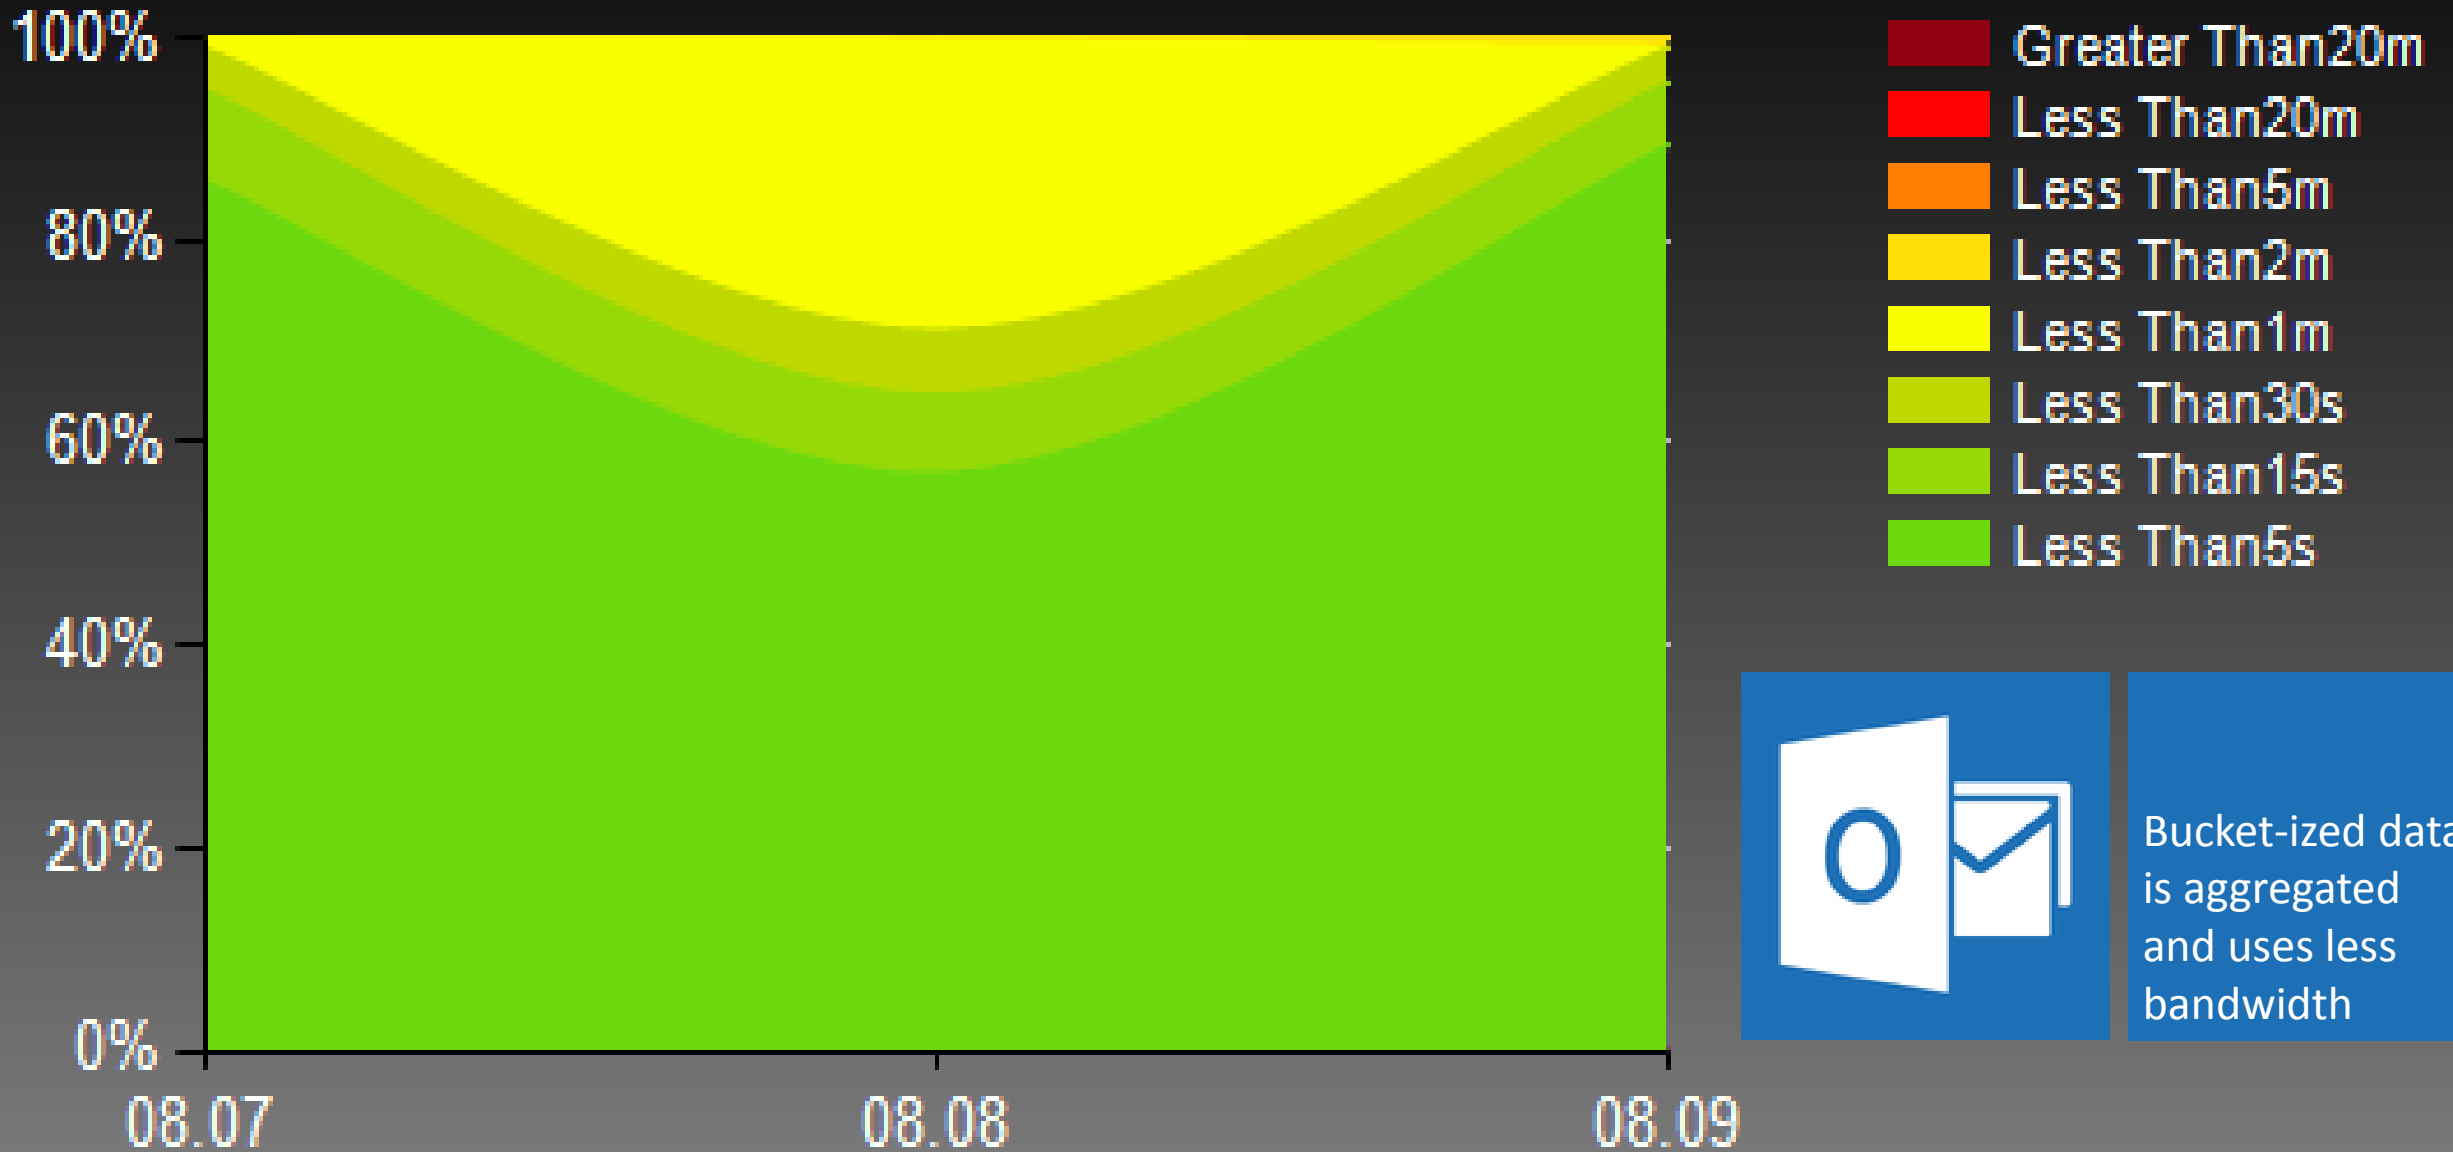
- Bandwidth

- Battery

- Disk



Initial Connect: HTTP



Data acquisition for clients

Filtering and aggregation at client

Be kind to the client

Pipeline to collect and process data

Make it easy


Staged Data Acquisition



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?



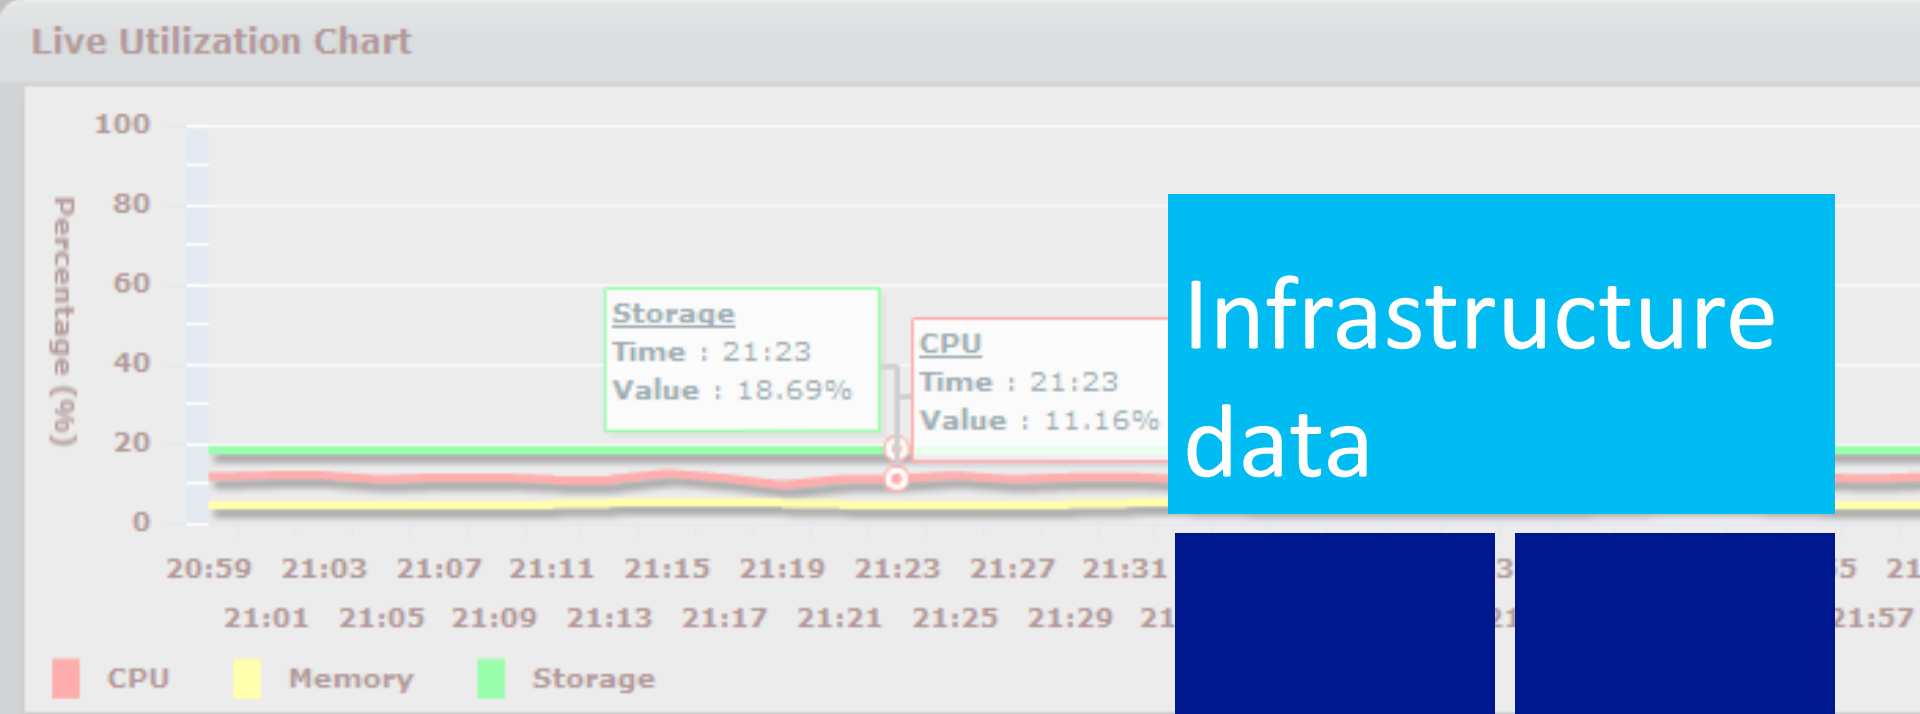
Select your data sources

Determine the data necessary to
answer your questions

Information

● **Online**

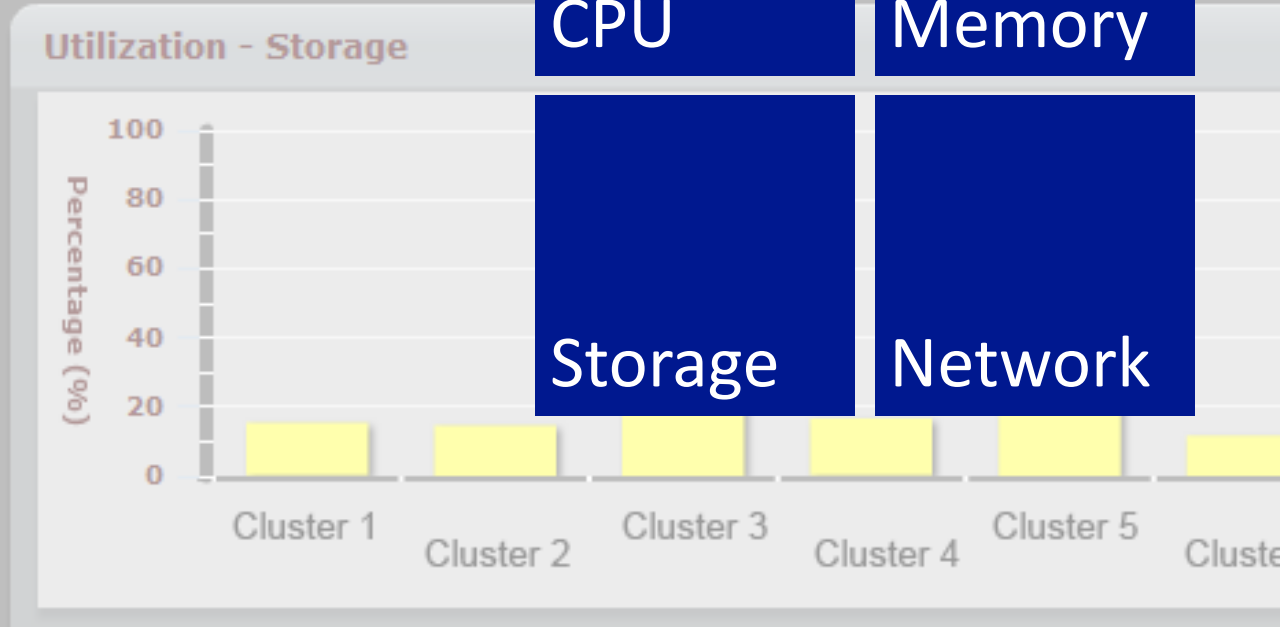
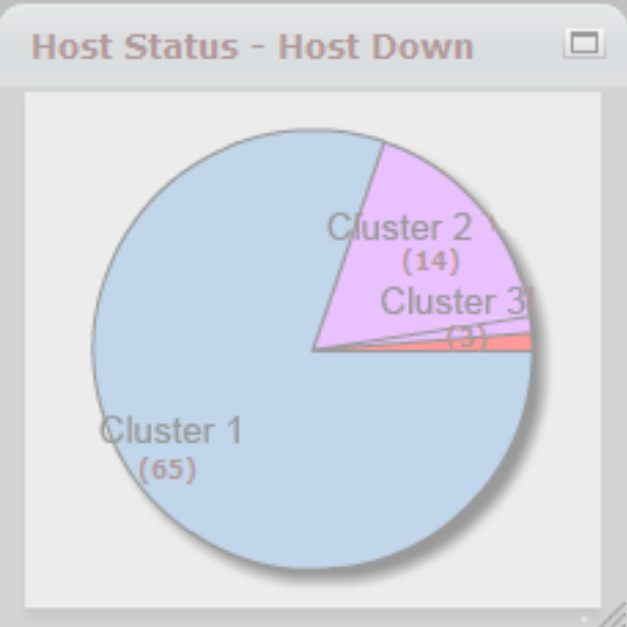
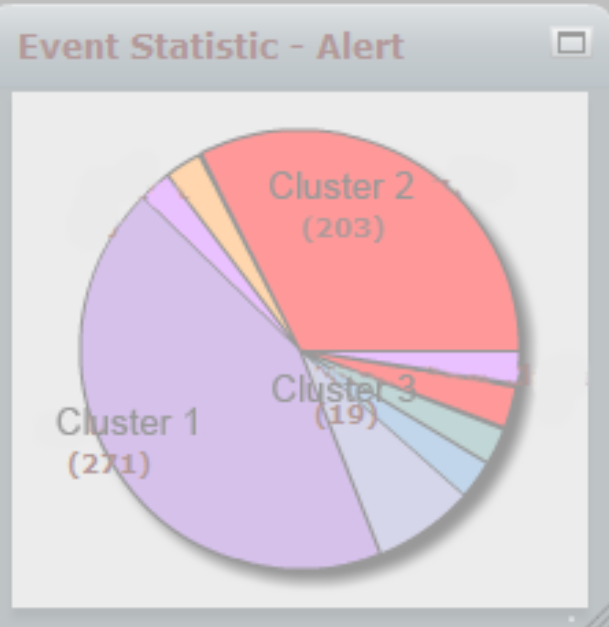
Total Cluster : 6
 Total Host : 45
 Total CPU : 502
 Total Memory : 1.78 TB
 Total Storage : 6.07 TB
 Processes : 37



Infrastructure data

CPU

Memory



Storage

Network

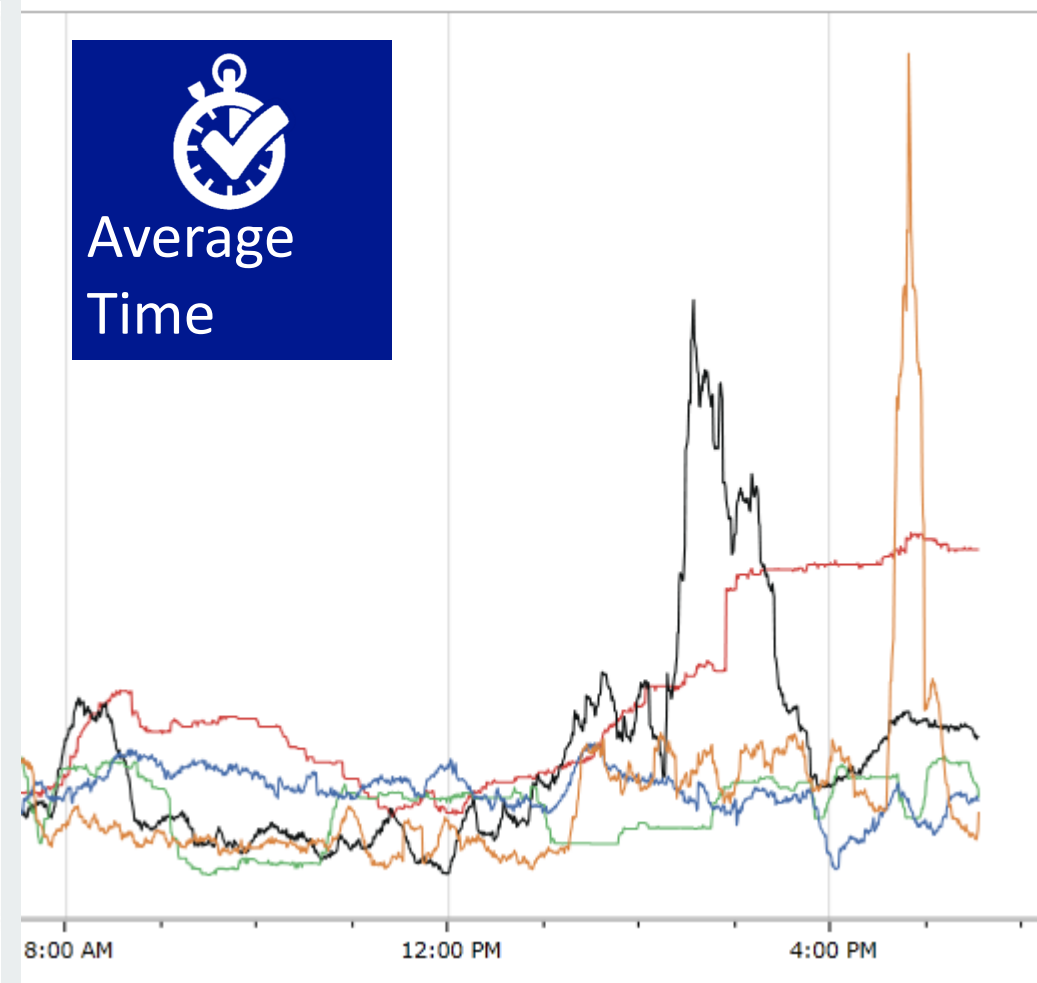
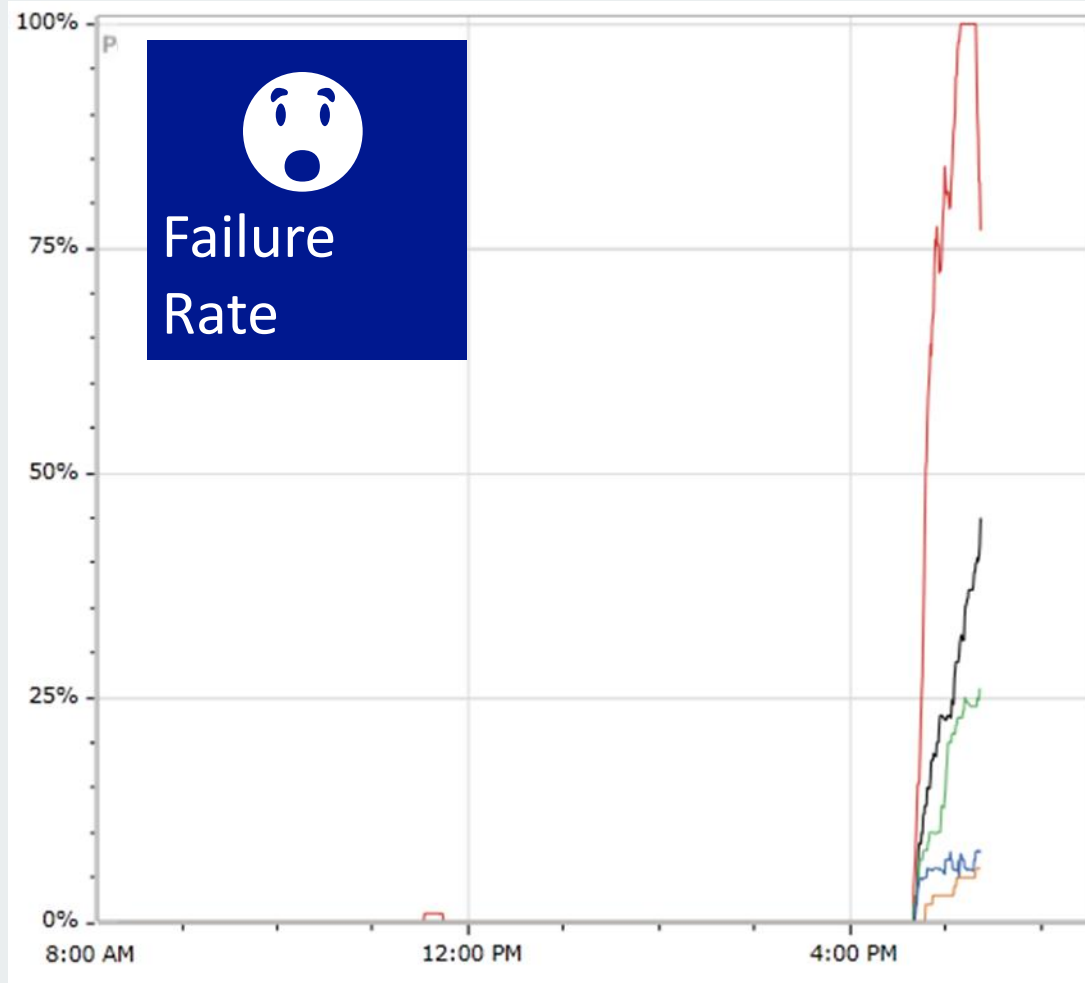
Application data



Office

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)

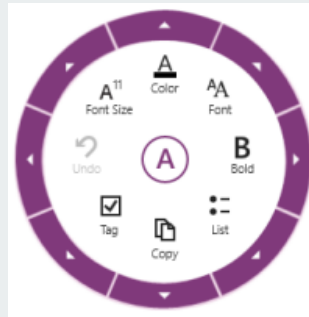
But it worked in Dogfood

Culprit was a old version add-in

Harden against this



Usage data



Client-side instrumentation

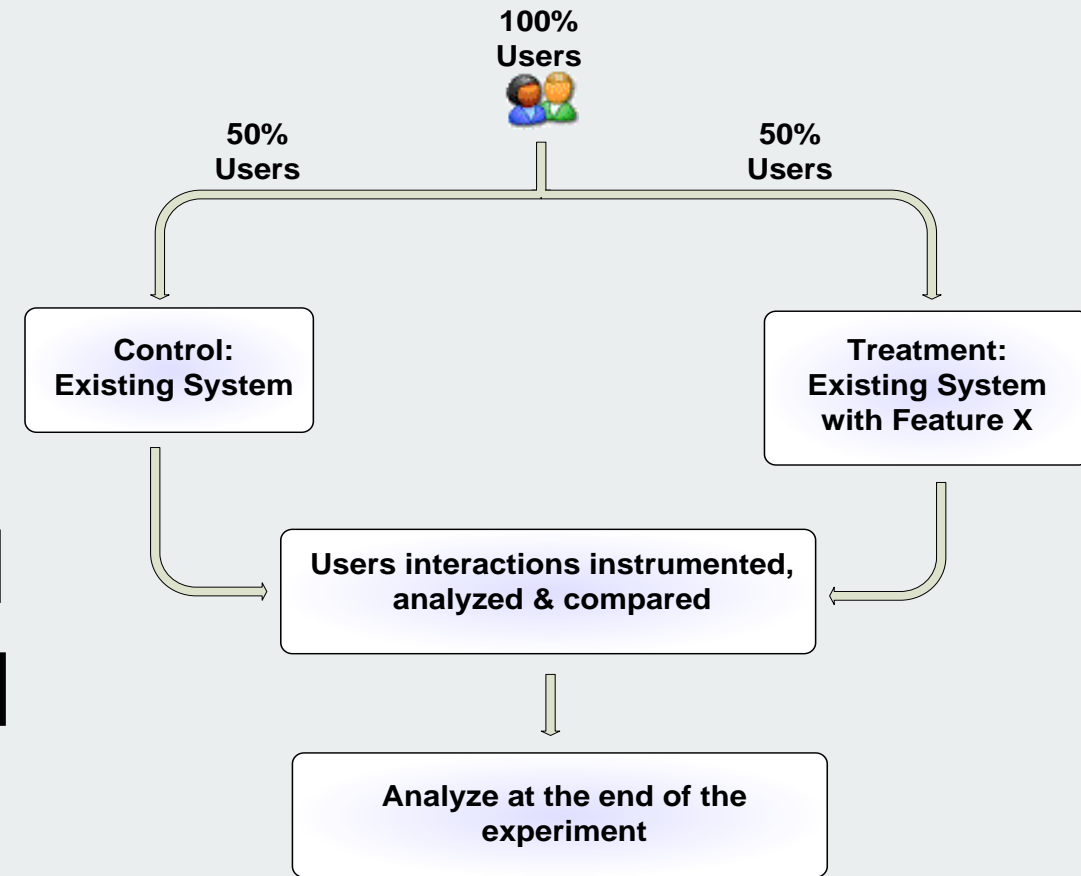
Proprietary

Javascript: clicks, hovers [web apps]

Get 1x1 GIF: Page Views [web apps]

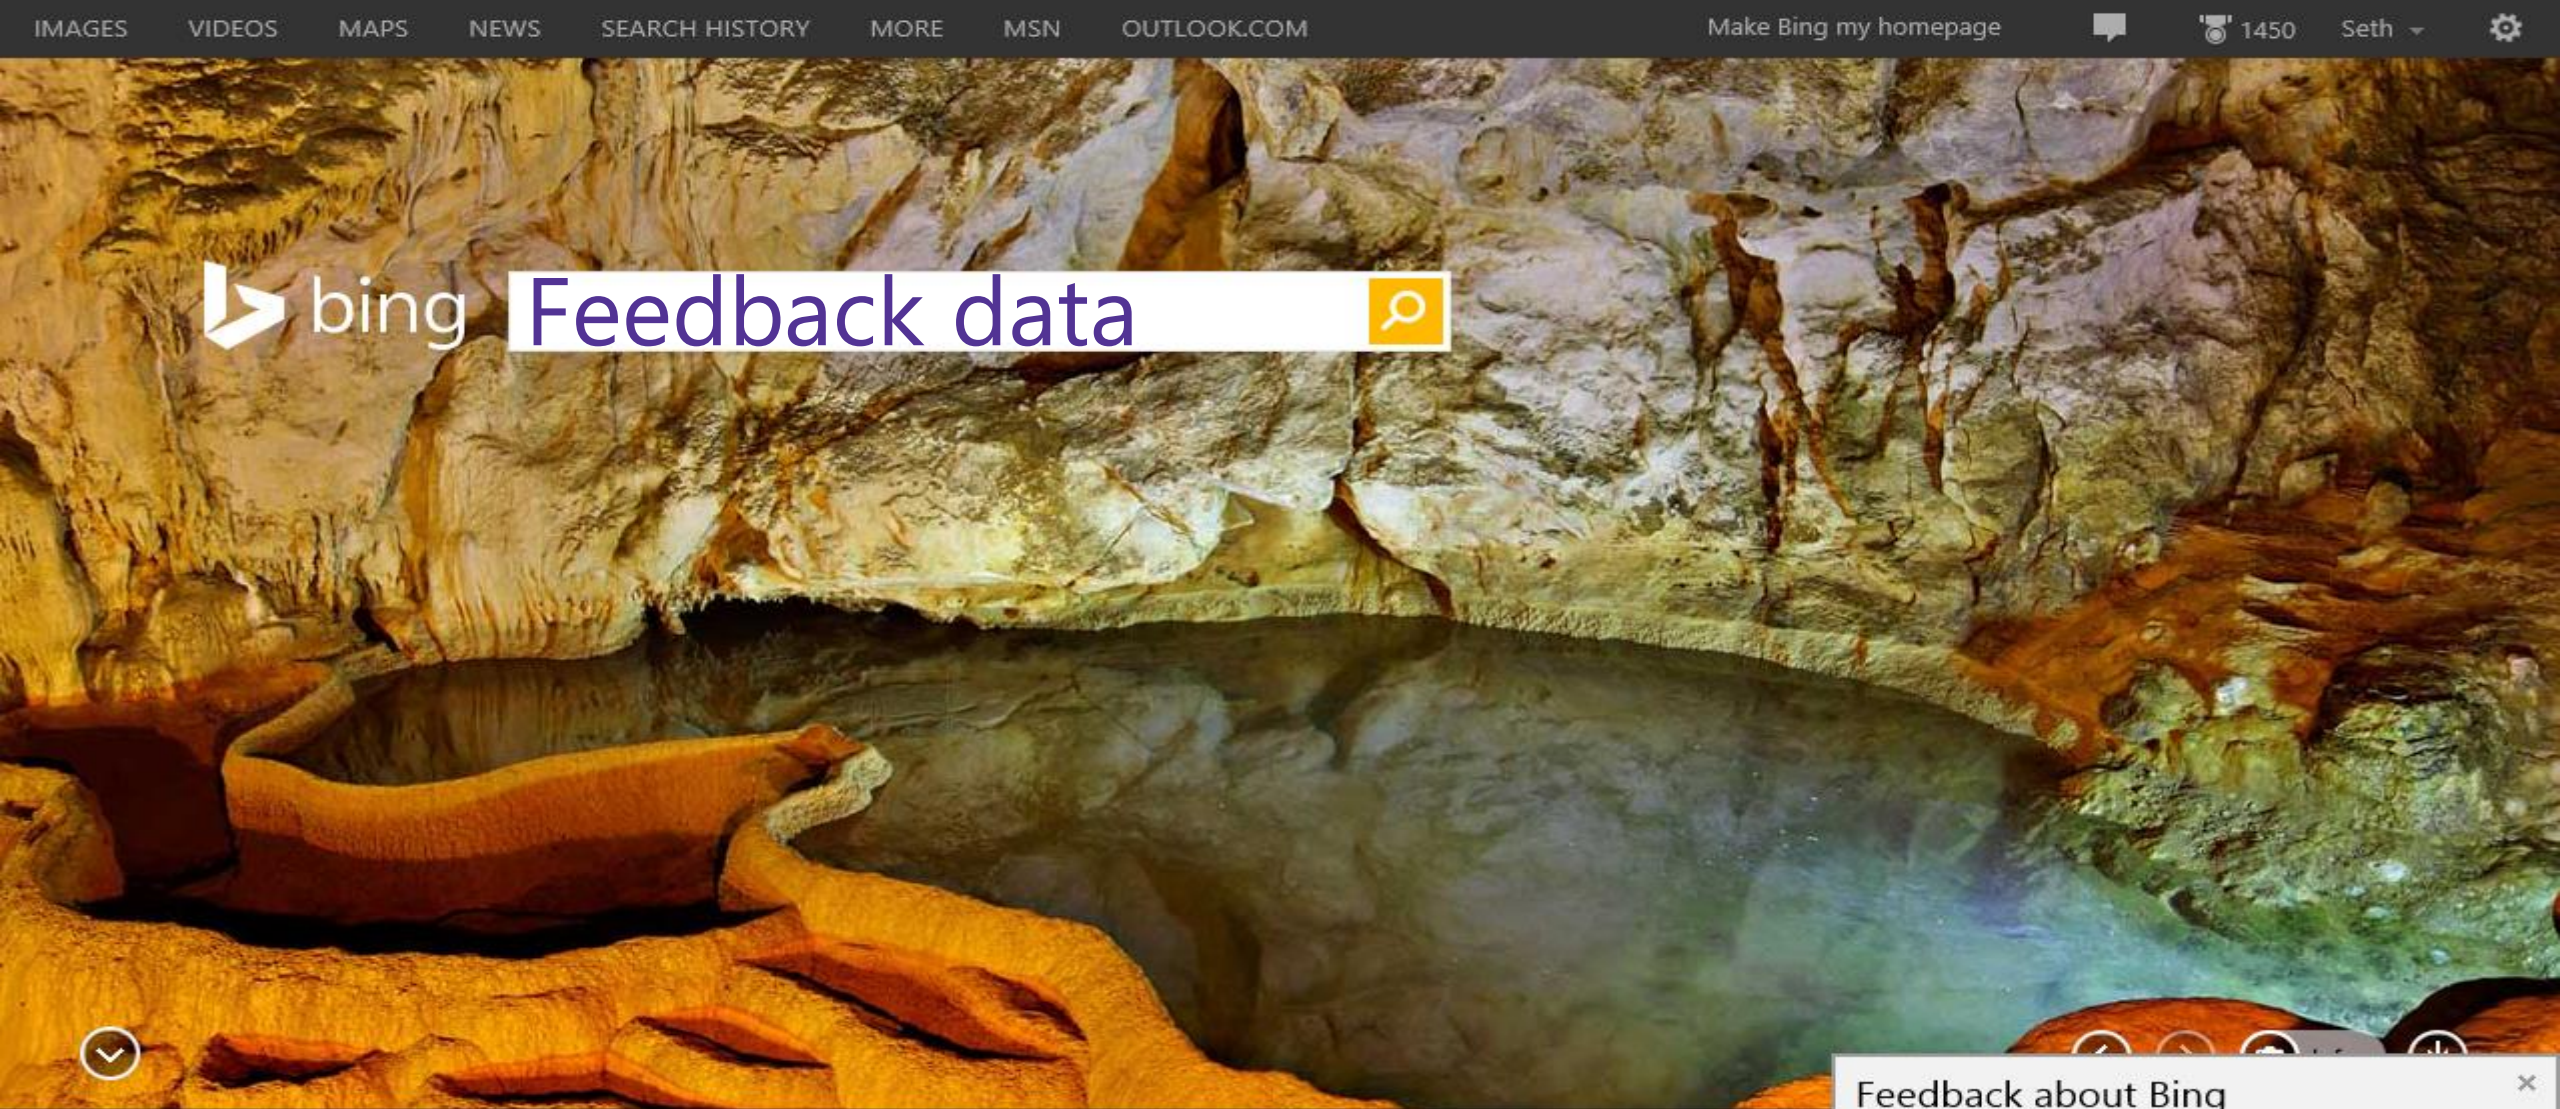
Combine into more complex scenarios

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**

The screenshot shows a Bing search results page for the query "Microsoft". At the top, there is a search bar with "Microsoft" entered and a "Go" button. Below the search bar, there are navigation links for "Web", "Finance", "Shopping", "News", and "Wikipedia". The main content area is divided into "RELATED SEARCHES" and "ALL RESULTS".

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 results

Microsoft Corporation Great Results

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 **Financial**

Customer service 800-642-7676 **24.50**

Search within microsoft.com ▲ +0.10 (0.40%)

US:MSFT Xbox

SHARE [Facebook](#) [Twitter](#) [Messenger](#) [Email](#)

Microsoft - Wikipedia, the free encyclopedia

History · Product divisions · Culture · Corporate affairs
Microsoft Corporation (NASDAQ: MSFT) is a public multinational corporation headquartered in Redmond, Washington, USA that develops, manufactures and supports a wide range of...

en.wikipedia.org/wiki/Microsoft · Cached page


Microsoft Support

The Microsoft Support home page is your support portal for Microsoft products. Download updates and find top issues, error messages, and troubleshooting tips for Windows XP ...

support.microsoft.com · Cached page

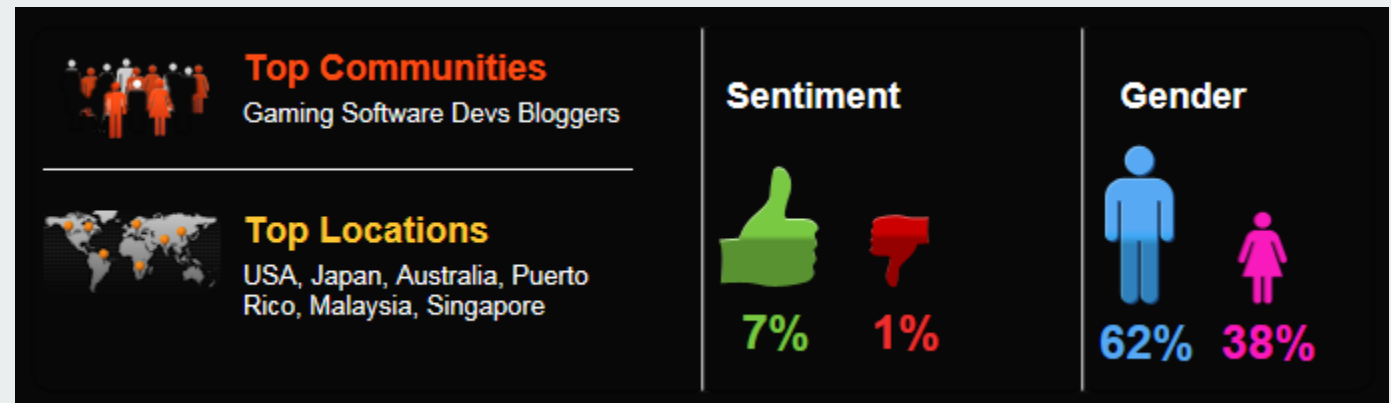
Feedback data also includes...

Customer Support Data

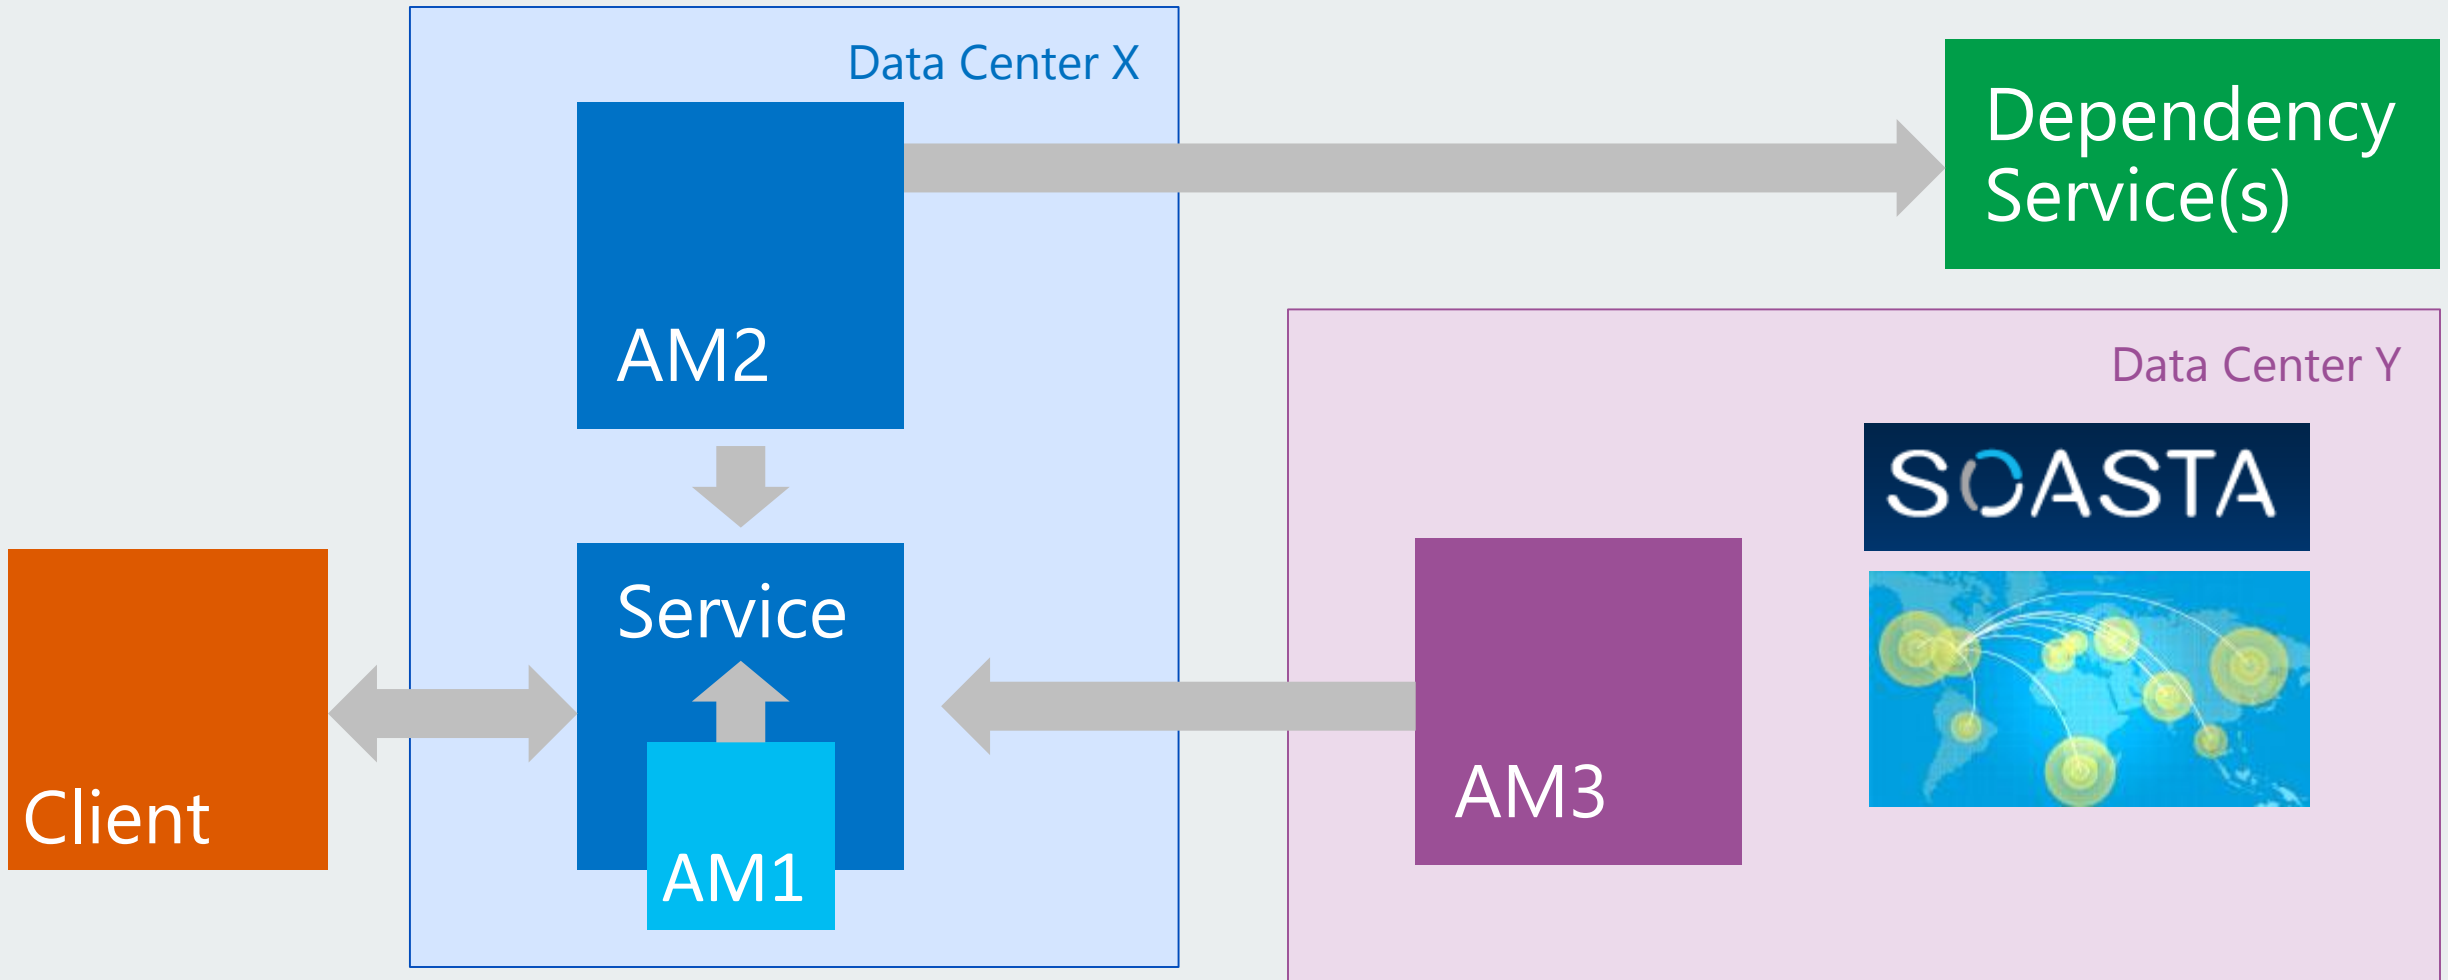


Service	Severity	Status
Lync Online	Sev1	➔ Service Degradation
Exchange Online	Sev2	🔍 Investigating


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





HDFS – Hadoop Distributed File System

Map-Reduce

→
ADDFCAECC

ADD
ECC

FCA

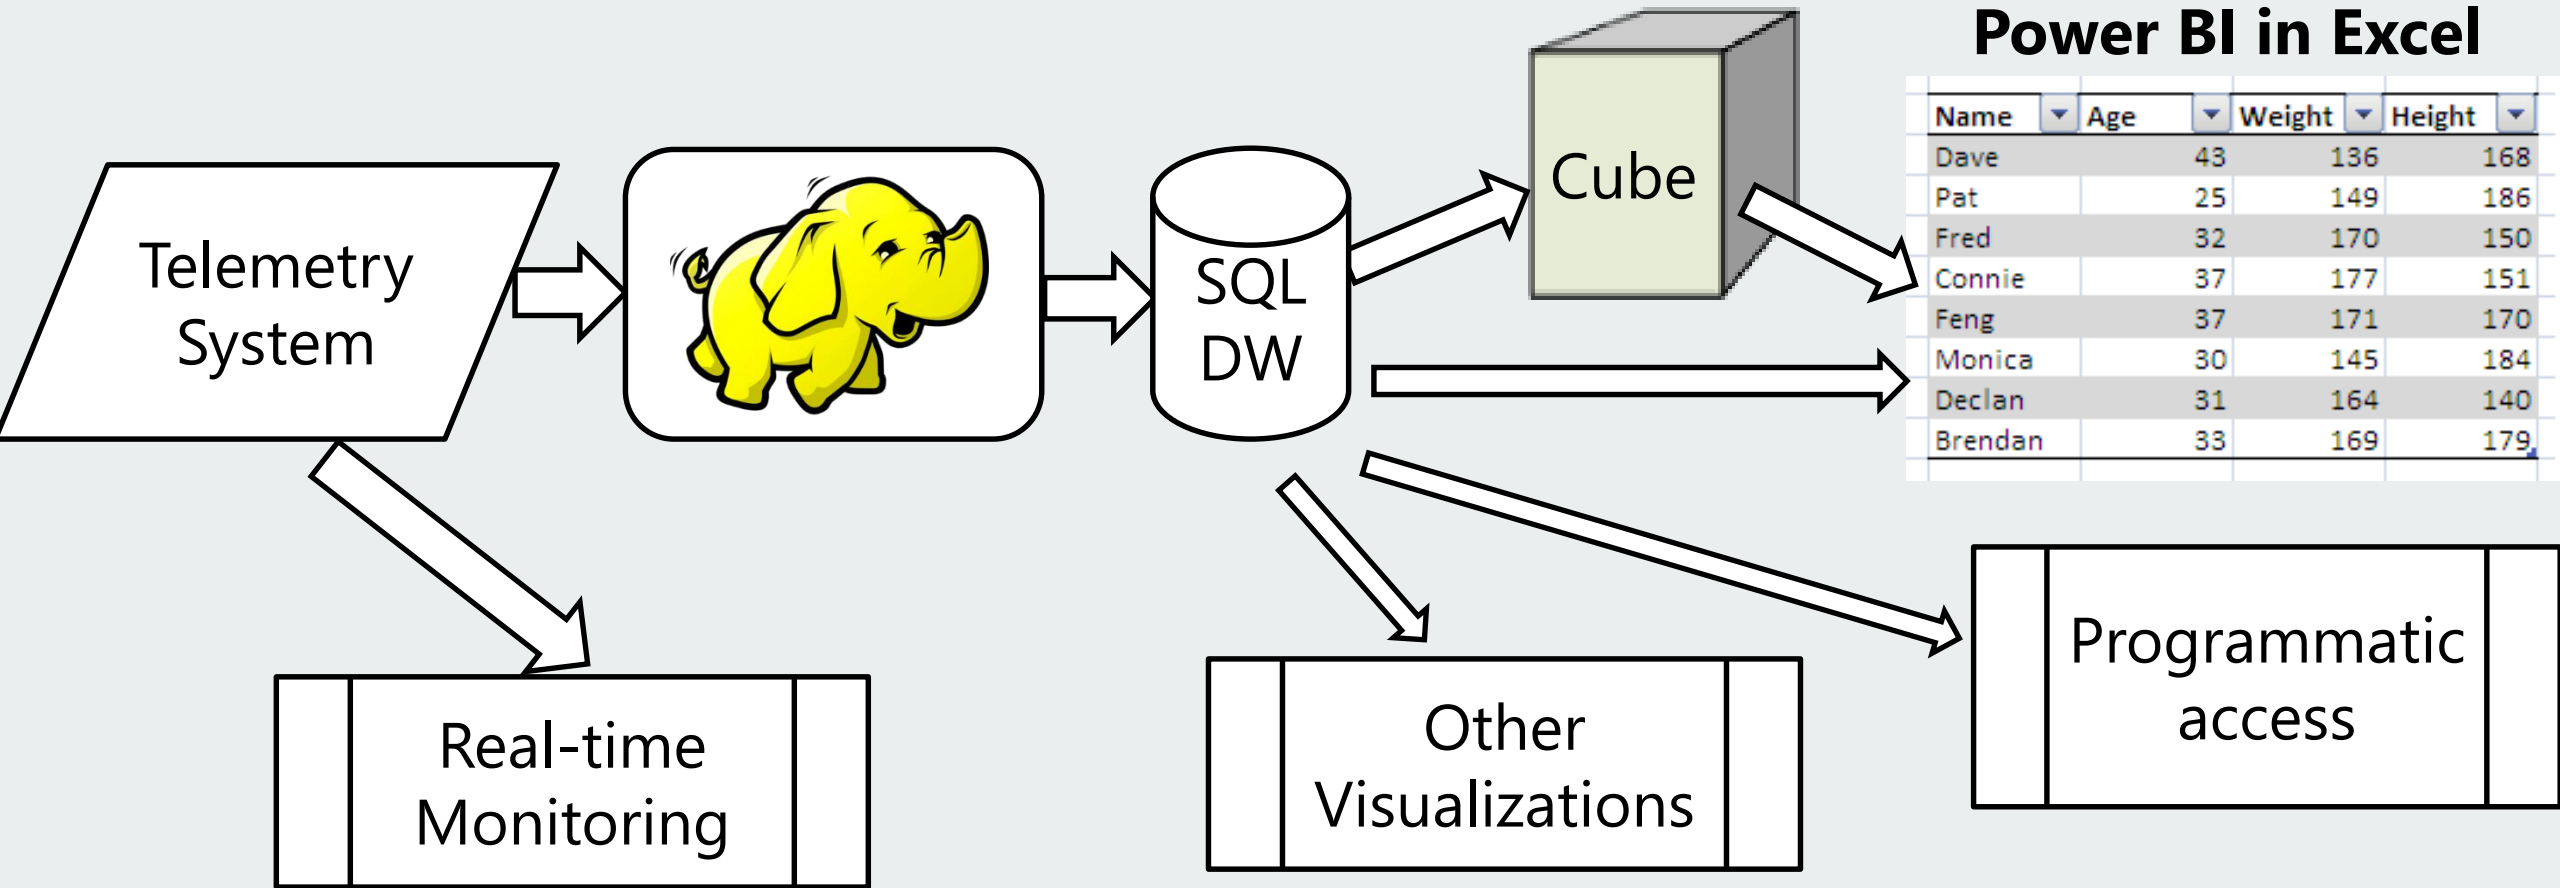


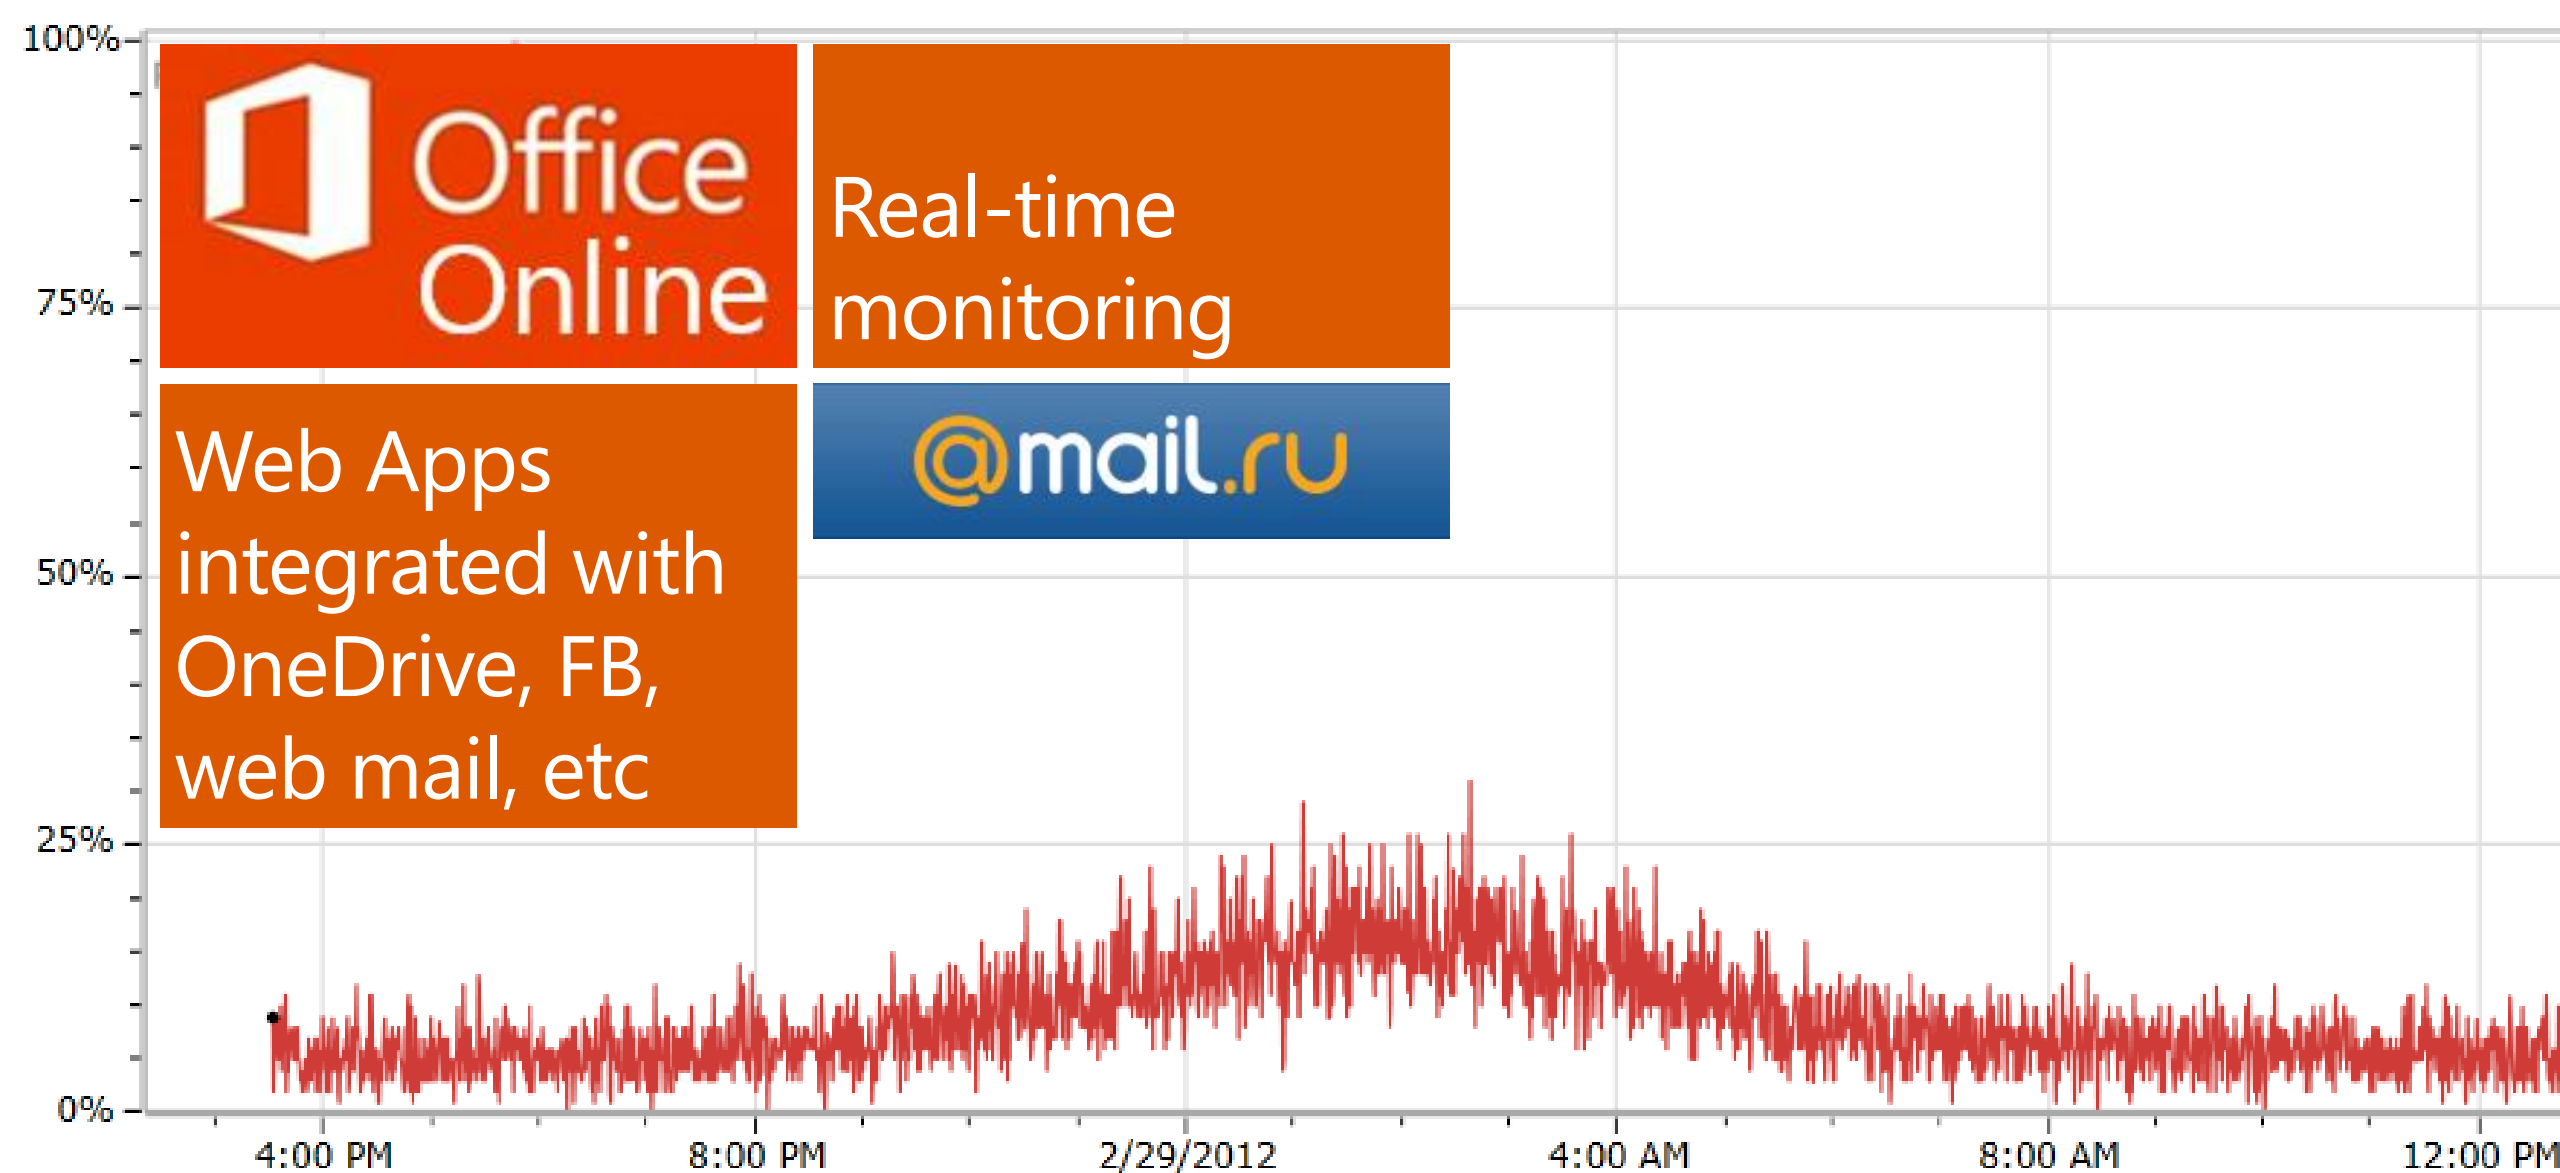
1xA
2xD
1xE
2xC

1xA
1xC
1xF

2xA
0xB
3xC
2xD
1xE
1xF
→




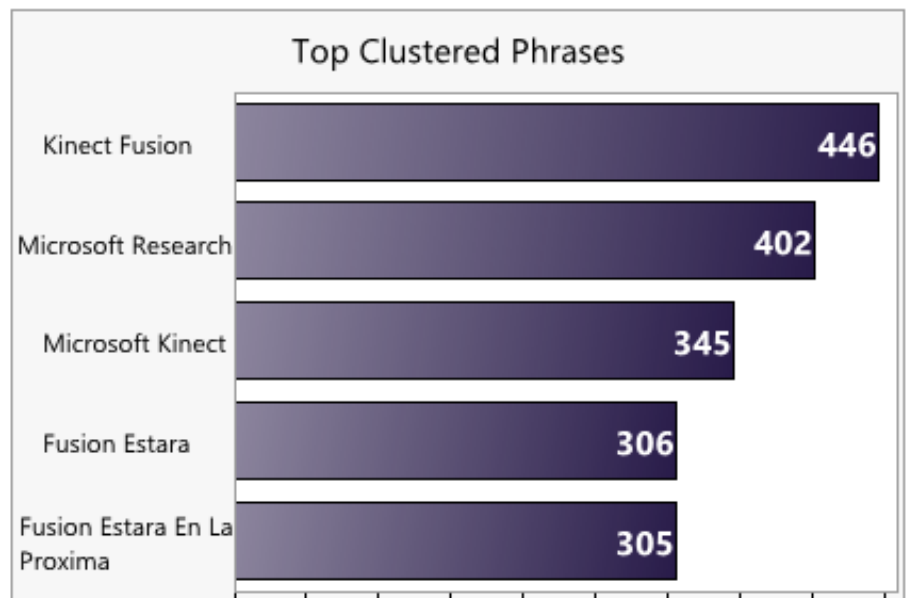


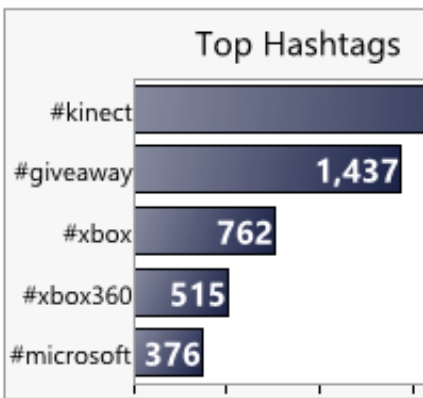



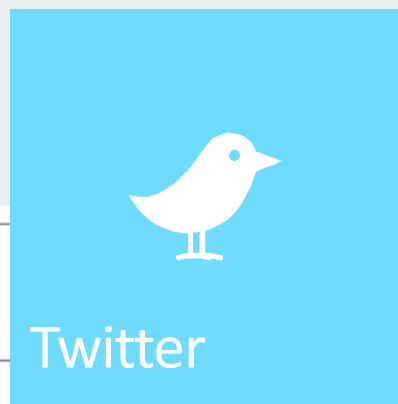
A common data flow





Pinned	Color	Metric	Instance	Aggregation	Value	DisplayedSum	DisplayedAvg	T
<input type="checkbox"/>	█	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?



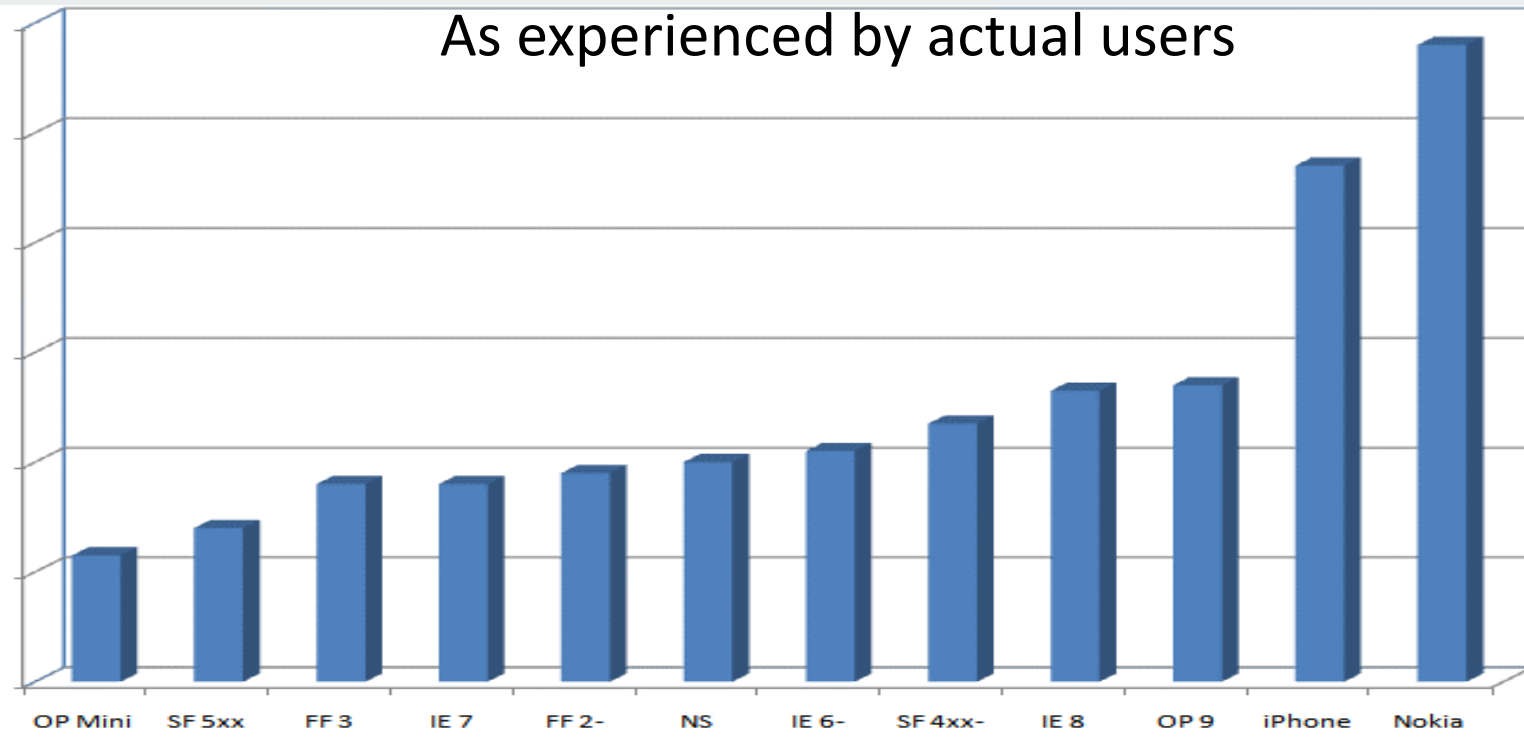
...sentiment analysis?



Get answers to your questions
and learn new questions

Outlook.com prioritizes performance

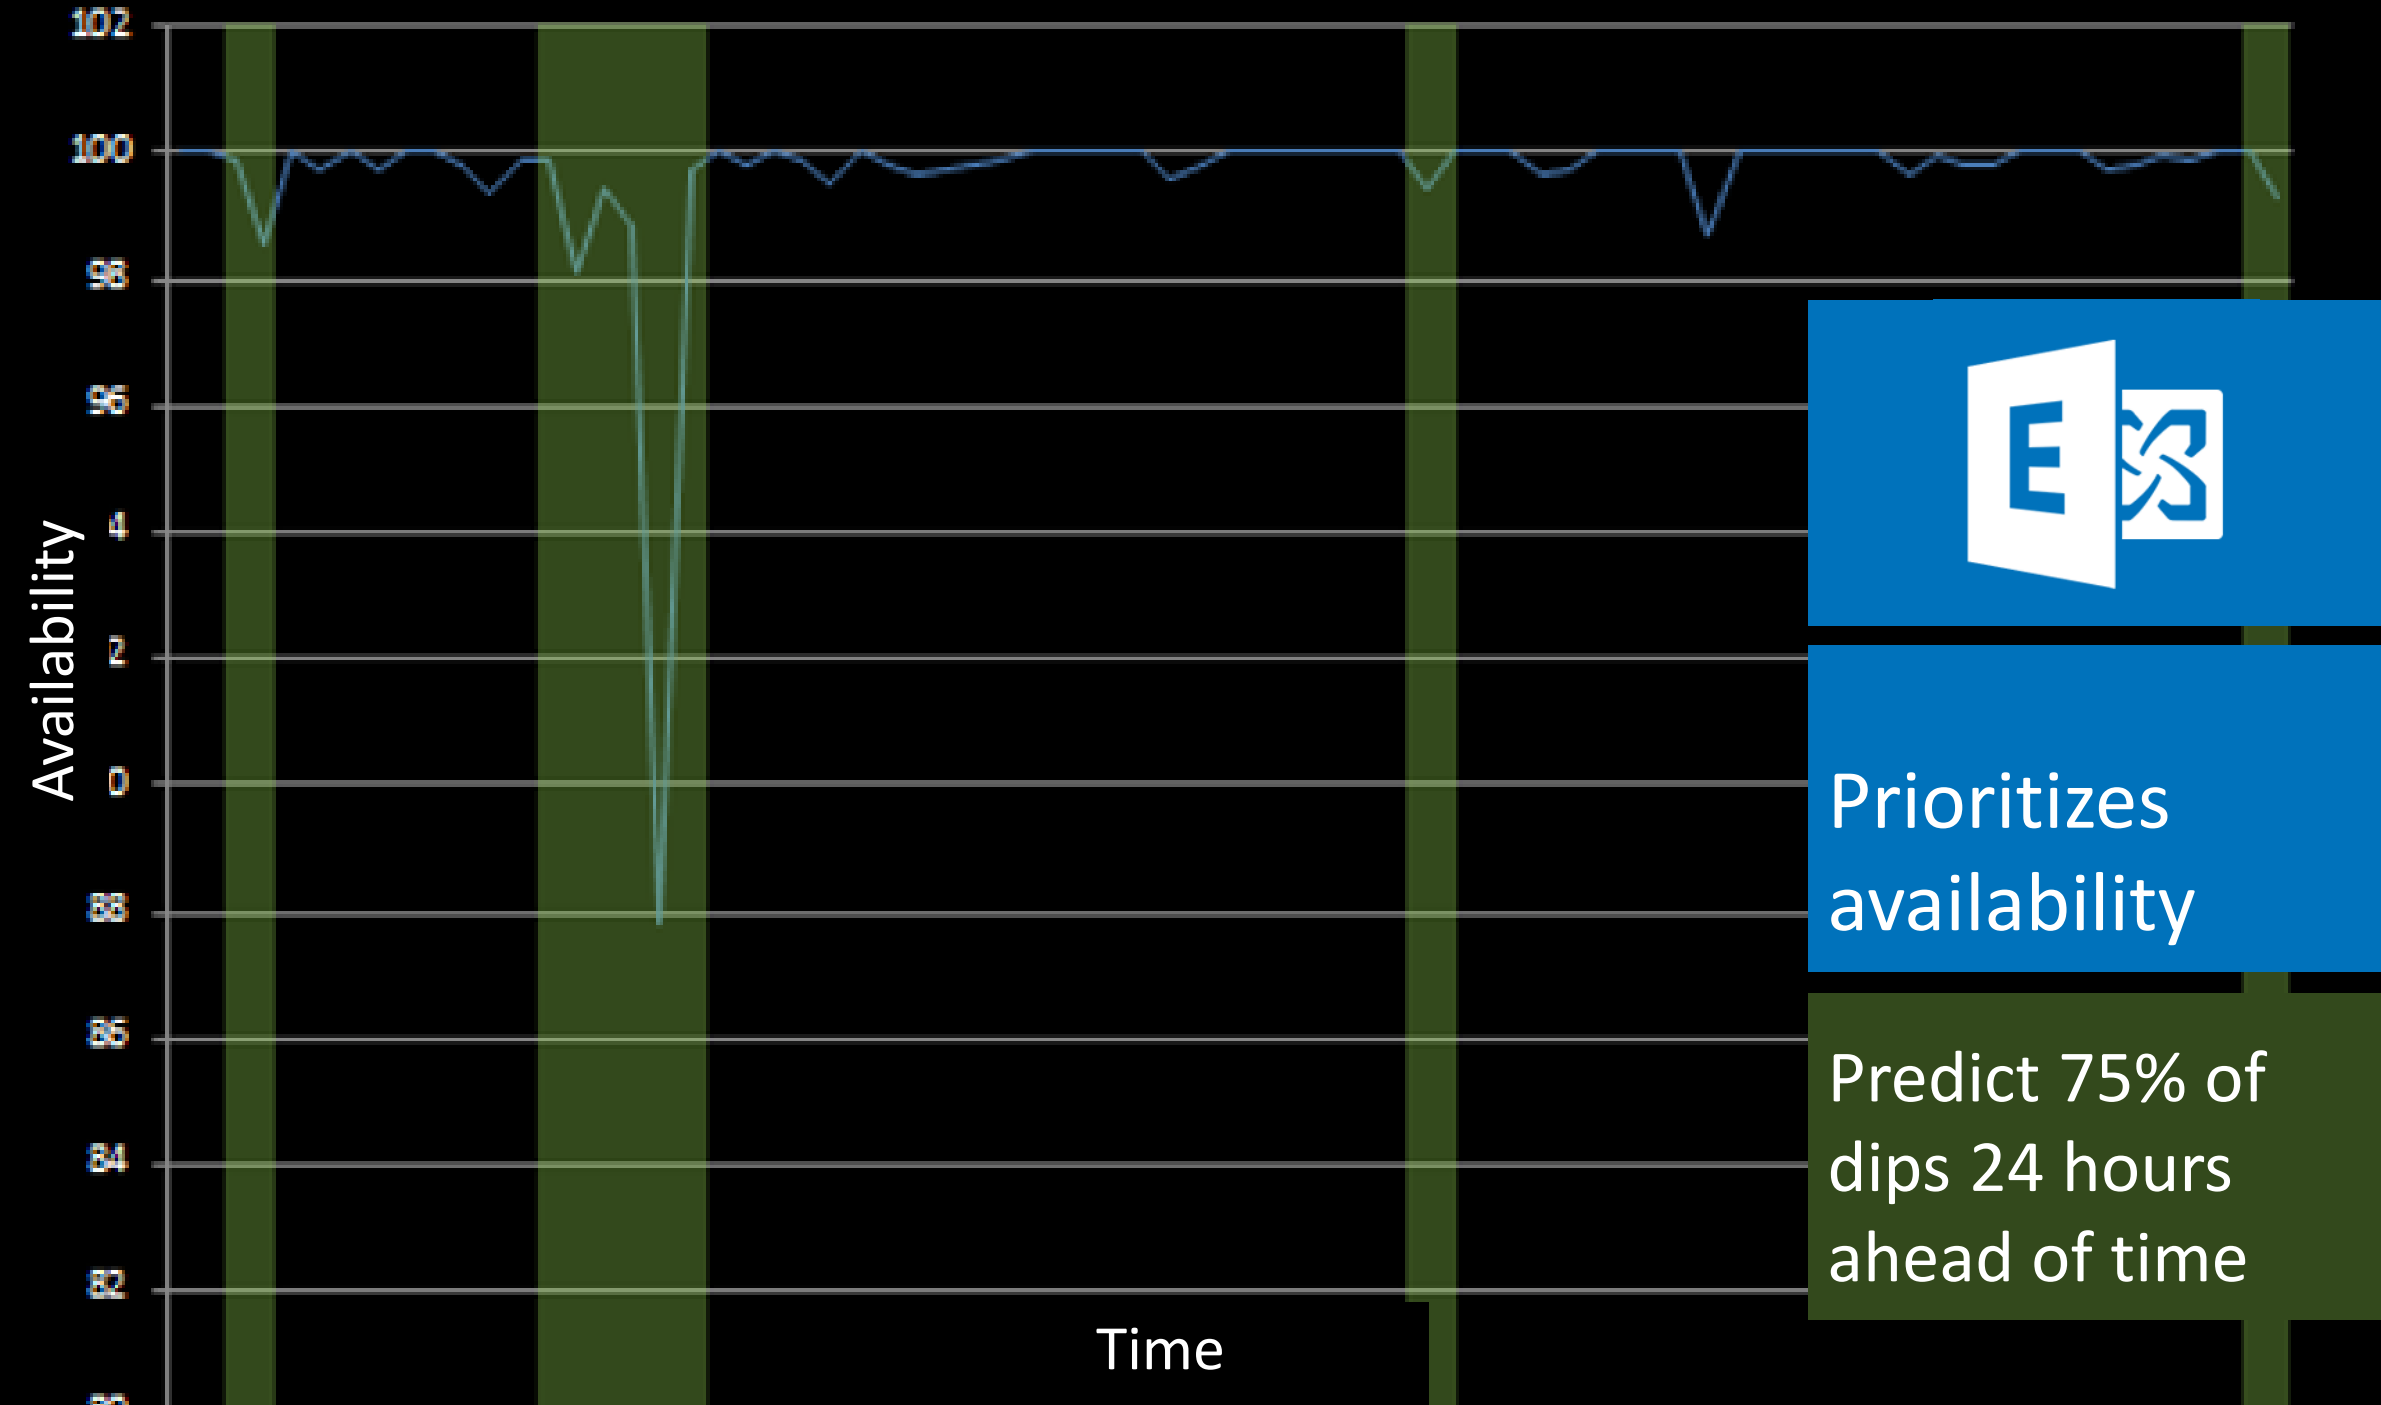
View Inbox – Page Load Time (PLT) by Browser
As experienced by actual users



JSI
JavaScript
Instrumentation

500 Million
measurements
per month



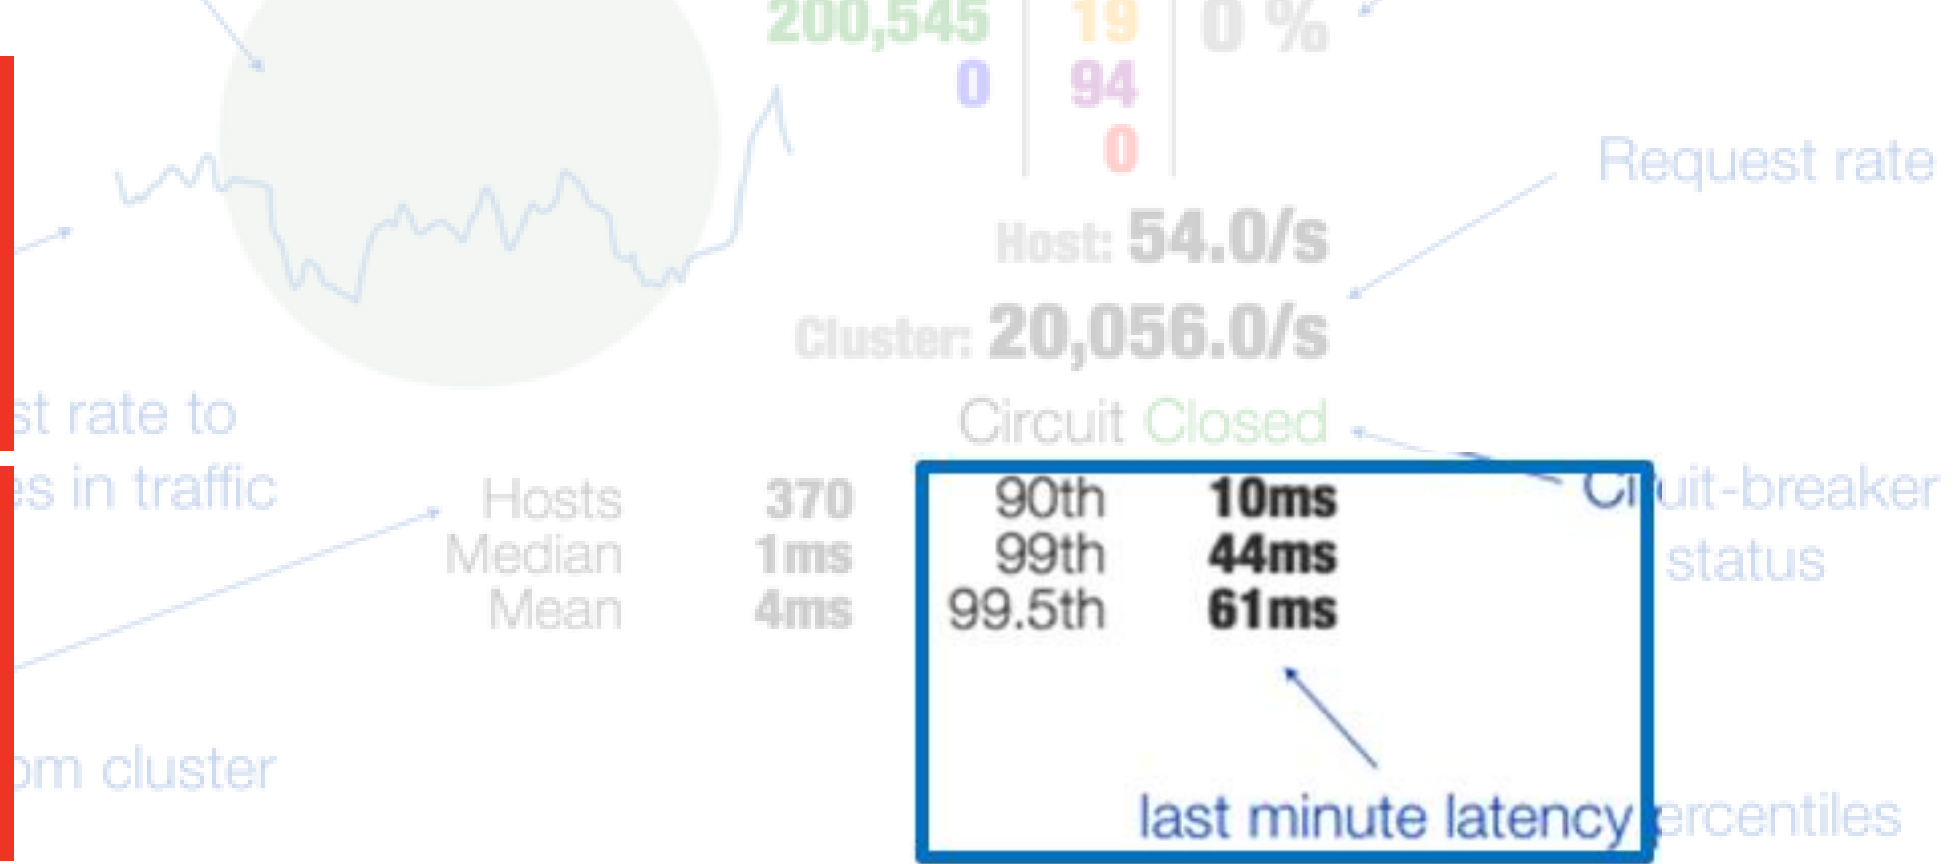


Prioritizes availability

Predict 75% of dips 24 hours ahead of 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

Step 1

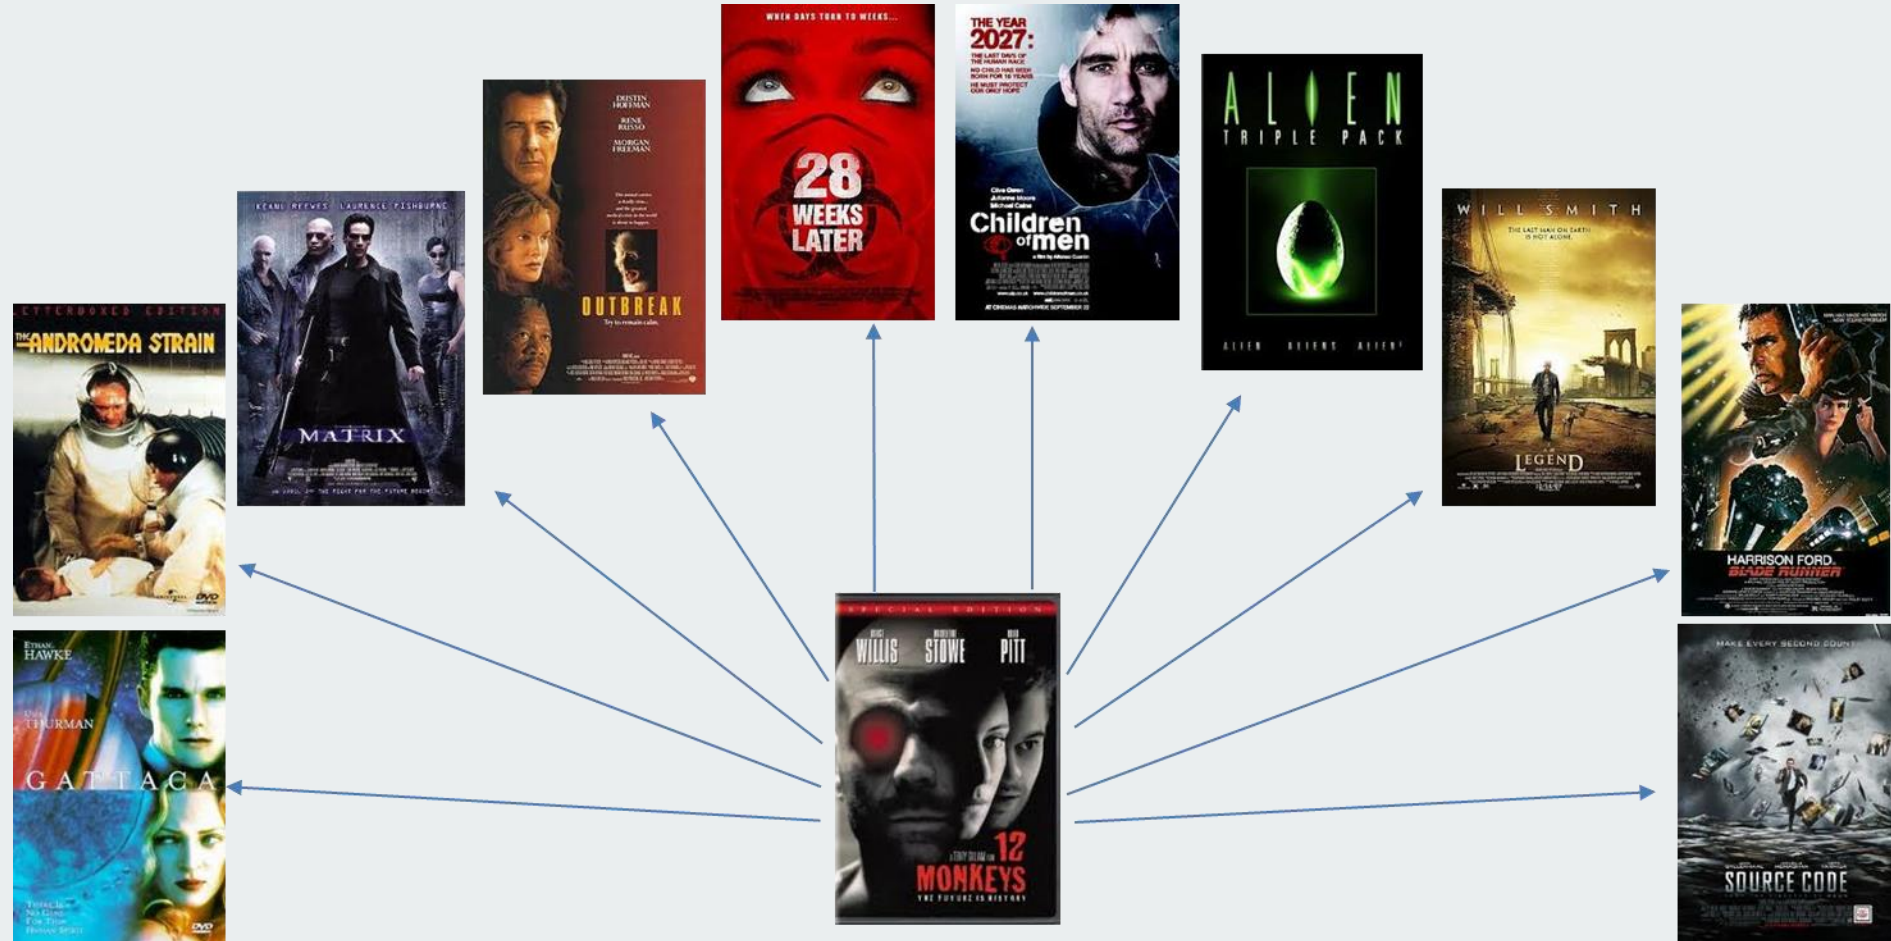
Step 2

Step 3


Step 4

Find new questions and repeat


- Xbox recommends
- Try algorithm
- Collect data
- Adjust algorithm
- Collect data
- Repeat



Find new questions and repeat

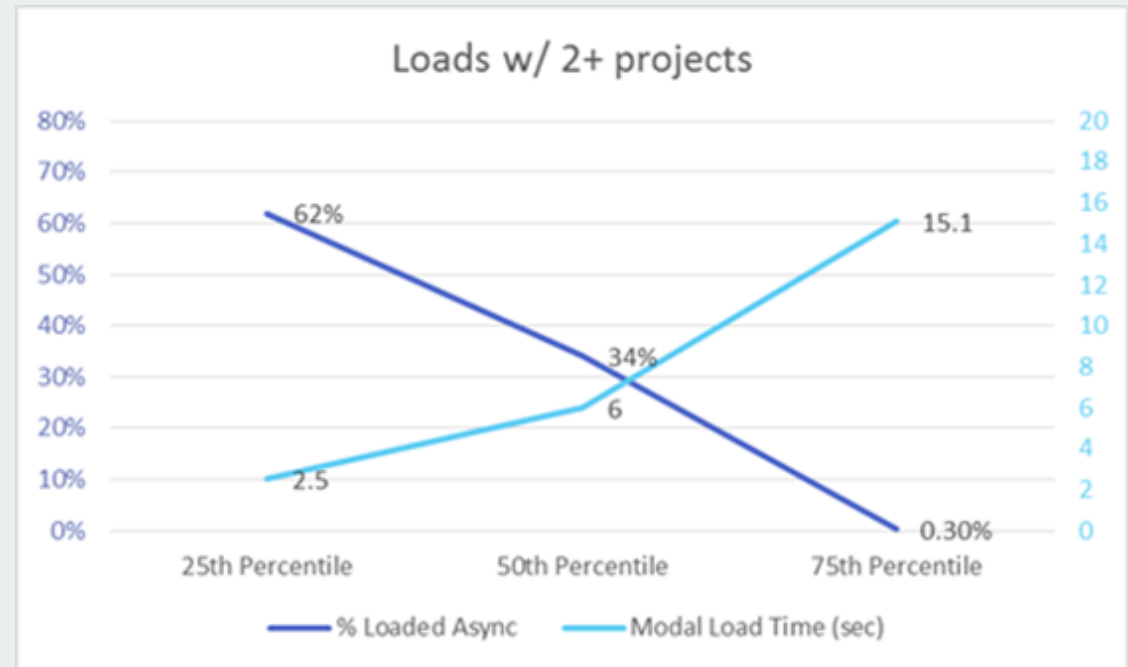


In Visual Studio 2012 we added asynchronous loading for *solutions*



We also added Telemetry

We can see how it works



Your turn –

Get answers; find more questions



Enjoy the rest of STPCon :-)

Get back to work



Implement your roadmap

Get 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?

Please fill out the survey



Create Your Roadmap to Data-Driven Quality

Session 101

Seth Eliot