



Stepping Up To Scale

Matt Davis
@dtauvdiodr
Platform Engineering, OpenX

Introduction

Matt

Musician

Maker

Databases

Systems

Synth Nerd



OpenX

Pasadena

Real Time Bidding

Advertising Exchanges

Publisher Monetization

Distributed Systems



*global presence
billions of daily ad impressions
multiple baremetal datacenters*



*EE with MDC repl
> 20TB, billions of keys
> 700 nodes deployed
8 different use cases*



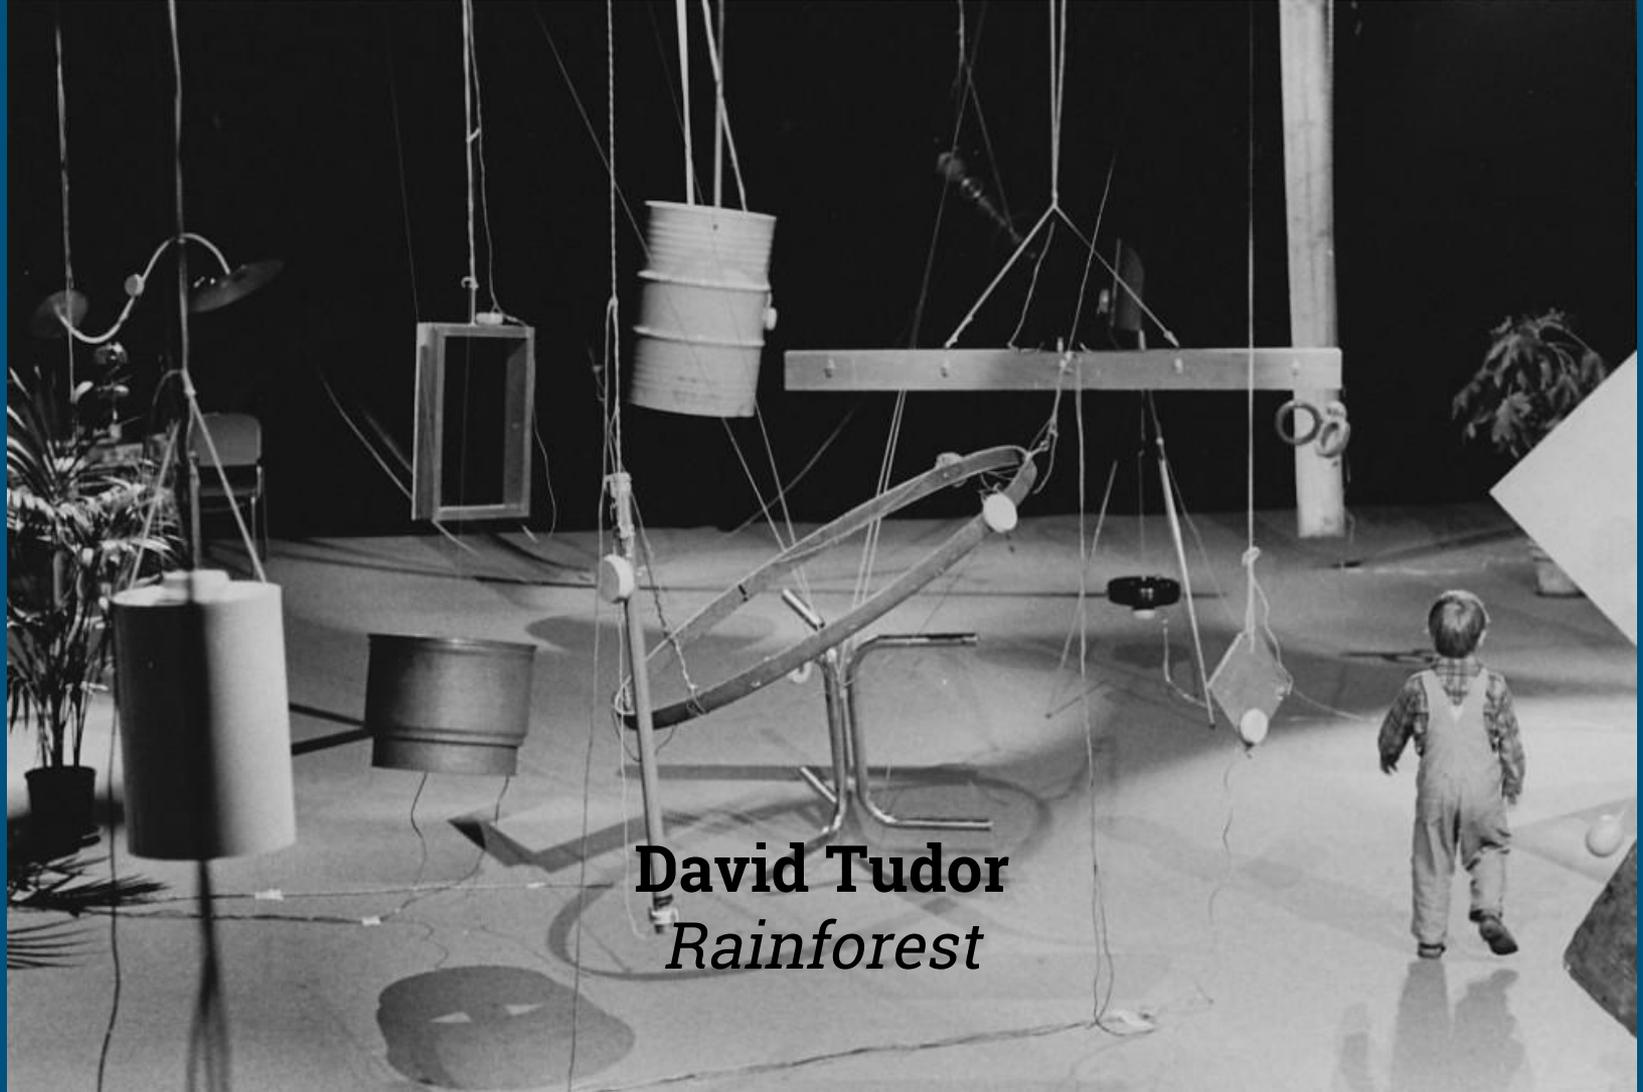
kafka



STORM

*> 10PB in multiple clusters
event and reporting jobs*





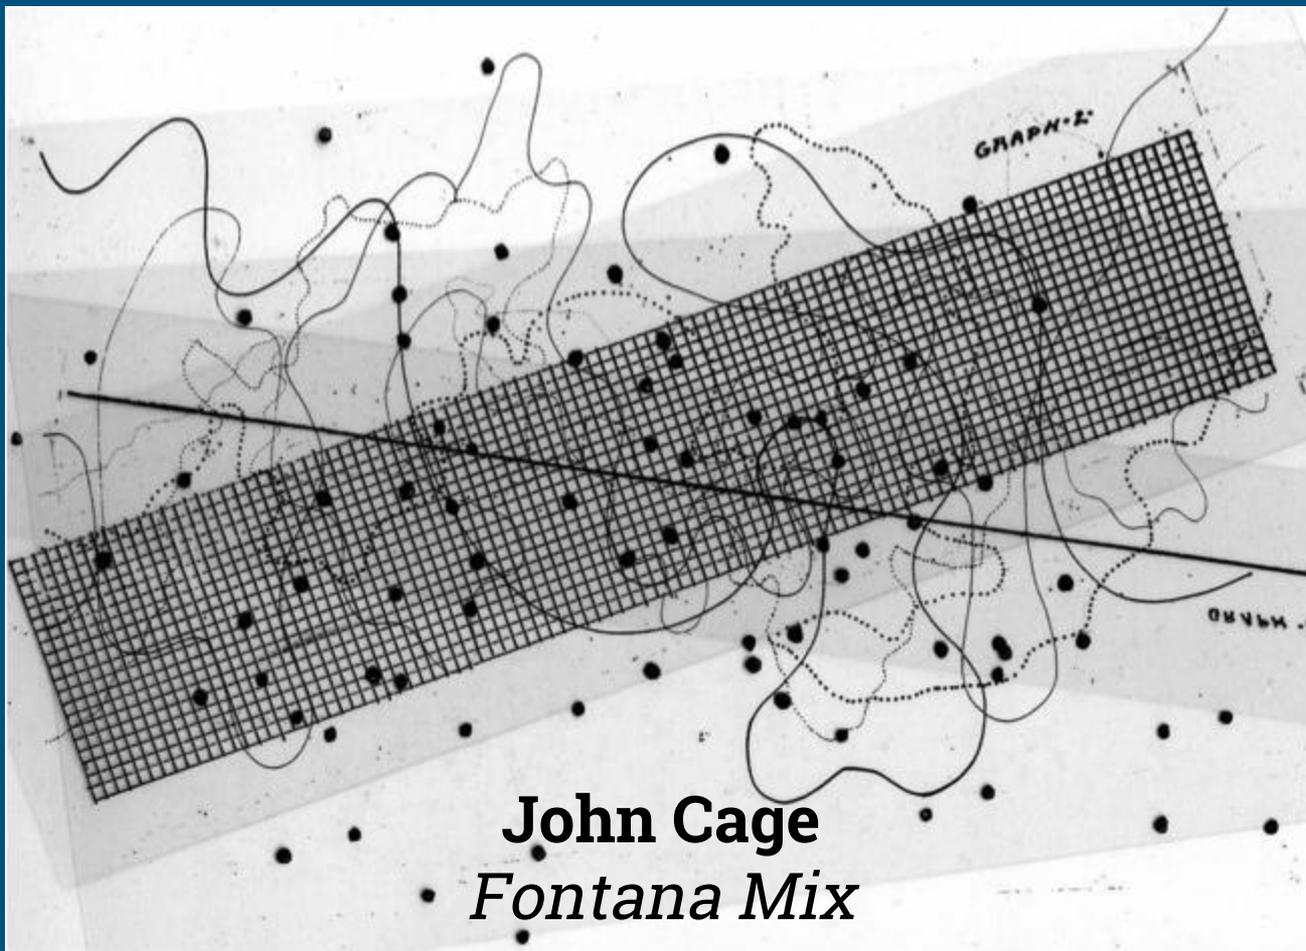
David Tudor
Rainforest

Indeterminacy as an Approach

John Cage built musical scores that were puzzles, indeterminate in respect to their performance.

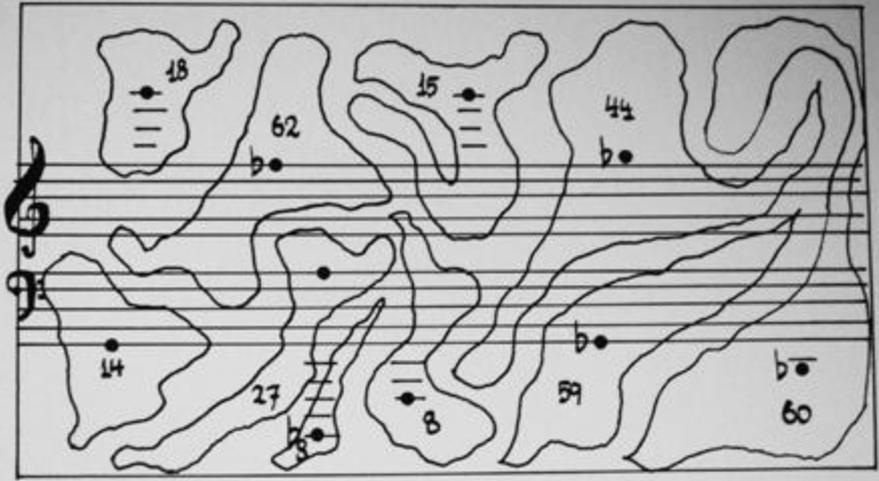
These scores do not prescribe what the sounds will be, or even their arrangement in time. They only describe the way events happen (or don't).

David Tudor has the job of mapping the indeterminate into something performable and repeatable.



John Cage
Fontana Mix

T



John Cage
Solo for Piano
(Concerto for Piano and Orchestra)

Indeterminacy in Music

placement of events in time

materials used

instruments involved

audience sound

external sound

Indeterminacy in People

In each of our lives, unexpected things happen to us *people* every day

Those events give life a structure we may not even notice

These same *people* use digital media, advertising is for them

Indeterminacy in Distributed Systems

Eventual Consistency

Actors don't know when an event will arrive

External: traffic spikes, market forces and seasonal bursts

Machine failure: down nodes (not so bad)

Machine failure: *partially* down nodes (worse)

Network partitions (expect them, always)

The Evolution of Ad Quality

An idea is born: *make online advertising bearable by providing quality*

Achieve minimum viable product with a multi-tenant, blended use case

A riak **bitcask** datastore (“TQ”) shared with a new **leveldb** one (“AQ”)

What Ad Quality does

Out of billions of bid requests per day, we identify ~500k unique creatives that are eligible for scanning

Intelligently scan partner websites with both image and audio/video recognition technology

“Creative ID” metadata is stored in a globally connected riak cluster

CRIDs are marked as block-able or not based on predefined rules

Real time impression stats for a CRID are fed through a Kafka/Storm pipeline into a MariaDB backend for analytical queries from various sources (e.g. UI)

Meanwhile the real time bidder process queries riak and denies the win to any marked as blocked

Surprise! It's GA?

The experiment has slyly become a product.

“Please build out this cluster by 5x by next month.”

So it's time to analyze the system...

*“**[Popular Automation Technology]** is really going to help with the sysadmin problem, they just don’t understand development.”*

- someone in a bar at a con

Principles of OpsDev

Don't be the Angry Sysadmin, but question everything

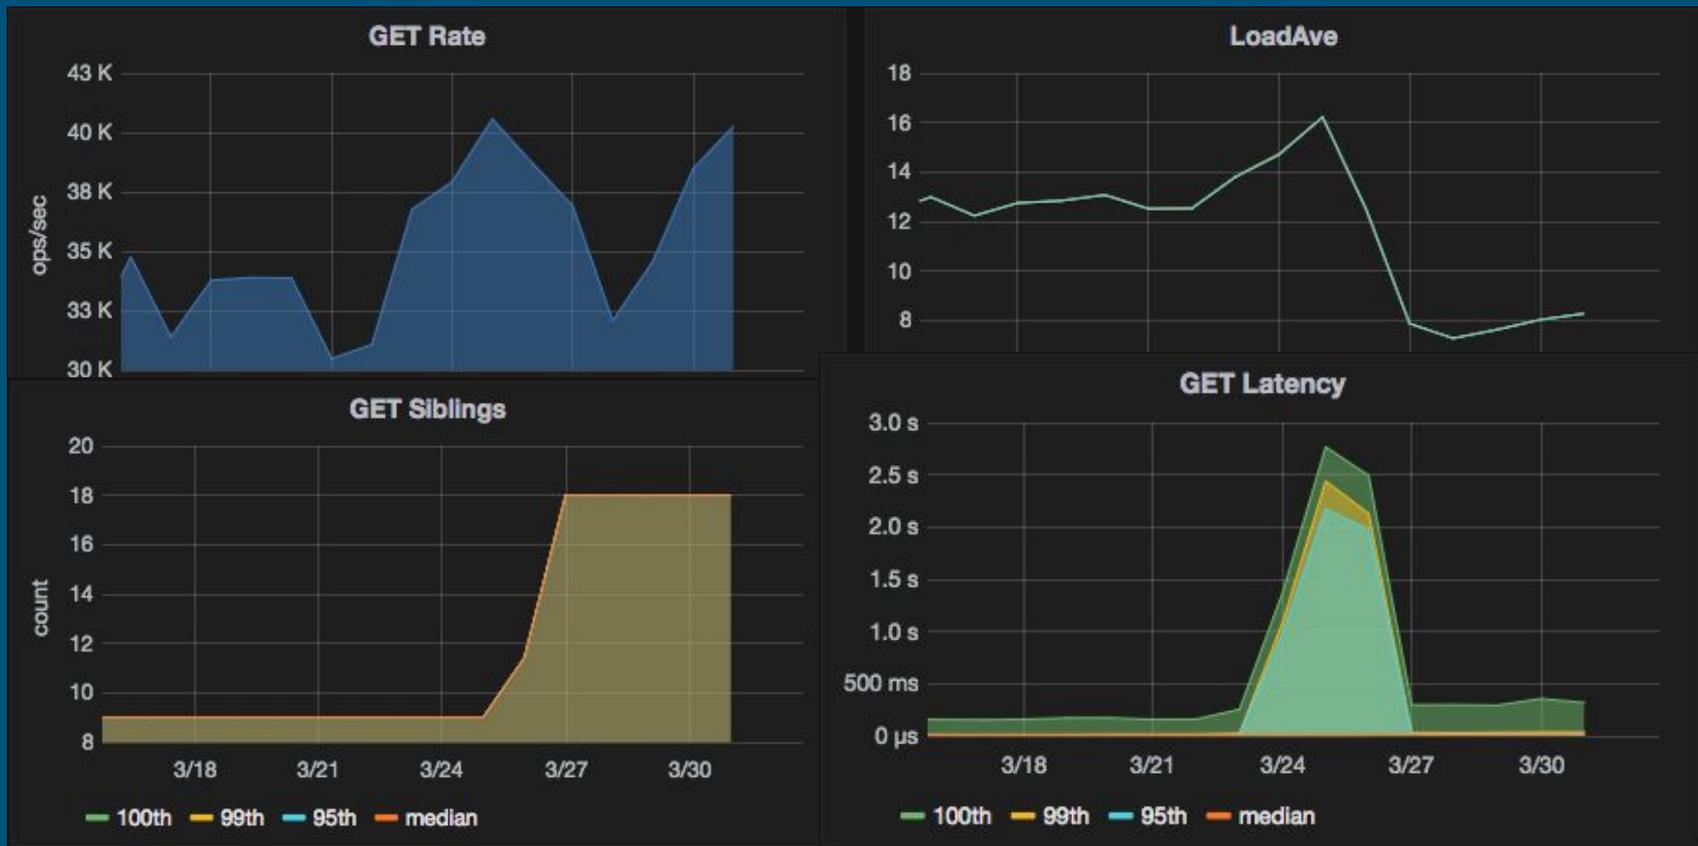
Reach back into dev, be present in their team, and educate

Internalize rhythms of the data

Provide the big picture, consider all angles

Know the flavors of *indeterminacy in the operation*

Indeterminacy strikes again: unexpected surges in creative traffic meant very high GET activity from the front-end delivery stack...



Time for some **OpsDev!**

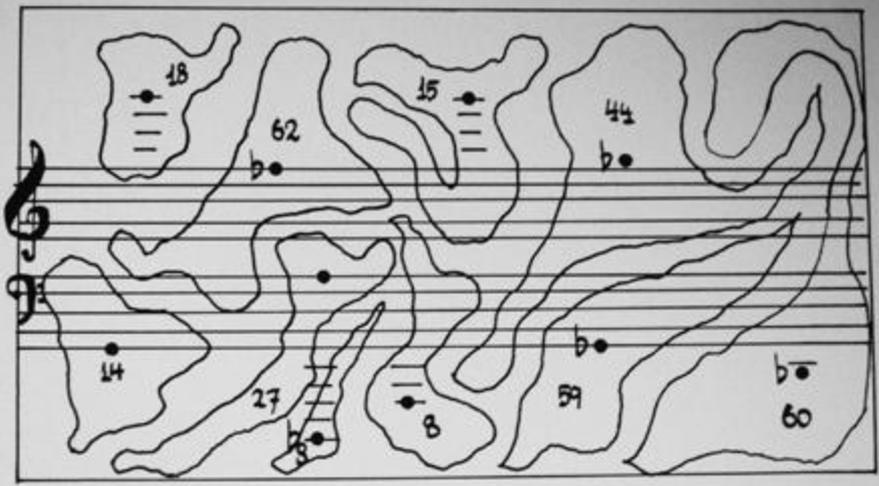
TechOps initiates a project with engineering to enable a cache layer so we can make improvements in the operation to allow for growth.



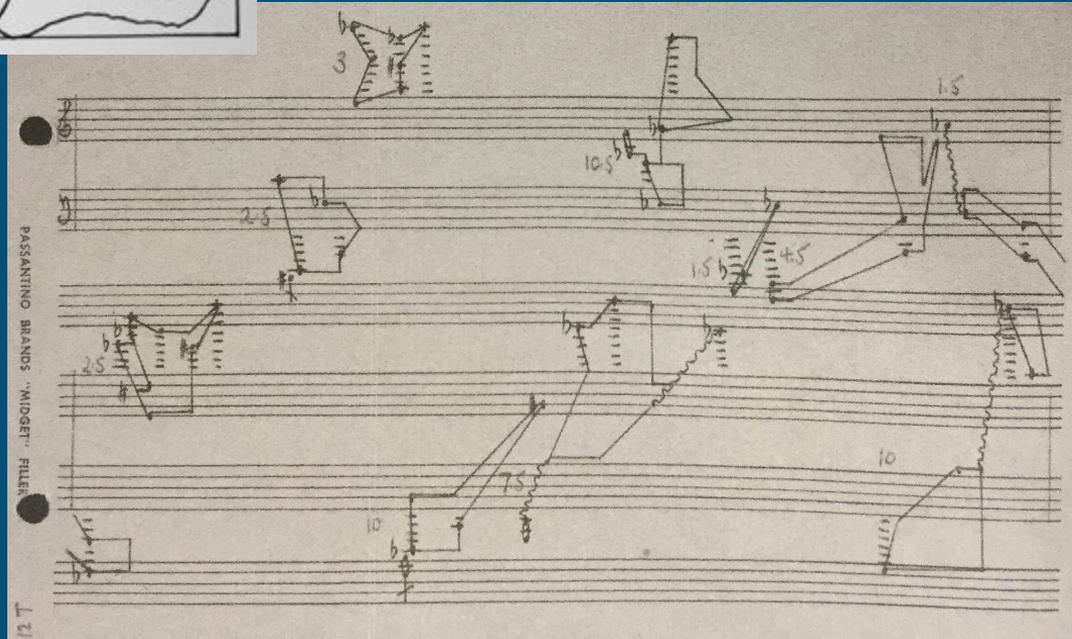
John Cage

Solo for Piano

*(Concerto for Piano and Orchestra)
the original score*



David Tudor
Reading of Solo for Piano
procedures for the "performable"
(and repeatable) version



Ops/Sec “performable” (repeatable) version

- Observe max possible rate per node (e.g. ~1500/sec)
 - Measure current trend of peak traffic in each datacenter (e.g. busiest at ~8000/sec)
 - Consider 30-50% CPU utilization
 - Add 85% capacity for Disaster Recovery
 - Assume future growth of 1.5x
 - Multiply by requested 5x scaling factor (the product request)



Minimal impact migration

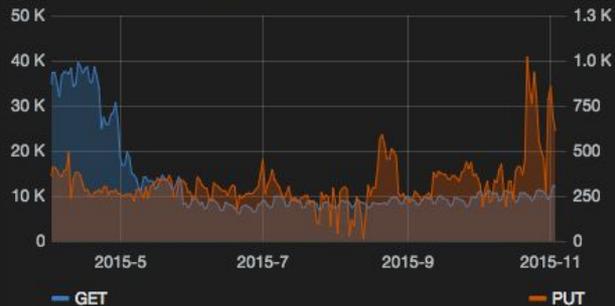
Using Riak KV EE with MDC replication, set up a one-way realtime connection from the old v1 cluster to the new (but equally sized) v2 cluster.

Briefly shut down the API layer, do a final fullsync, and switch configs.

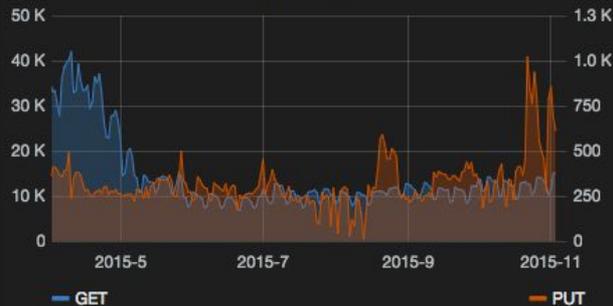
Now build out the cluster to expand scanning to cover all participating bidders, making it *really* GA.

Transition: Ops/Sec

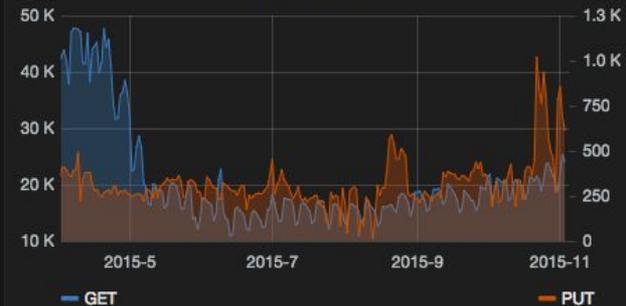
CA Riak TQ Ops/Sec



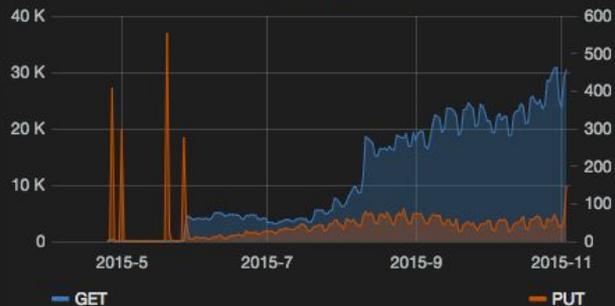
LC Riak TQ Ops/Sec



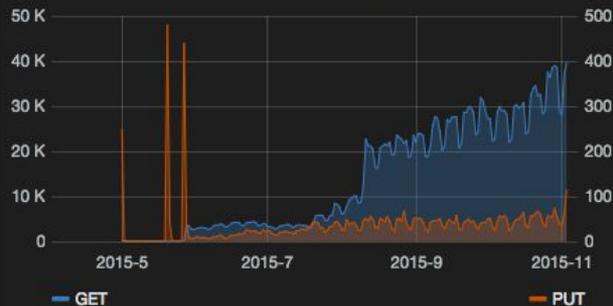
XV Riak TQ Ops/Sec



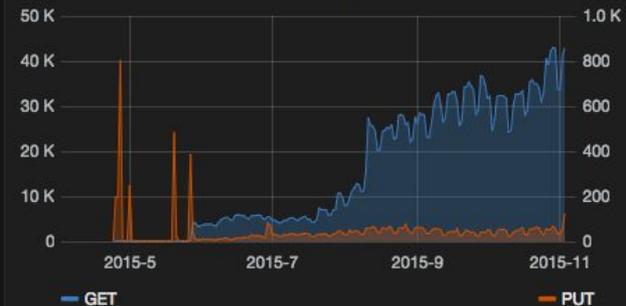
CA Riak AQ Ops/Sec



LC Riak AQ Ops/Sec



XV Riak AQ Ops/Sec

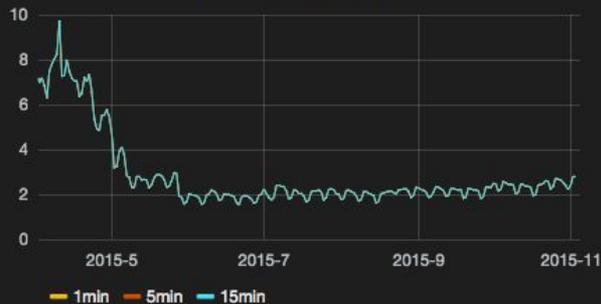


Transition: CPU

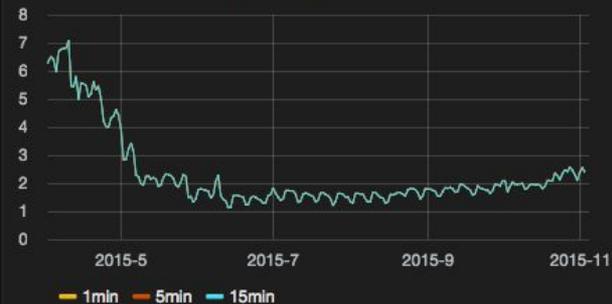
CA Riak TQ LoadAve



LC Riak TQ LoadAve



XV Riak TQ LoadAve



CA Riak AQ LoadAve



LC Riak AQ LoadAve



XV Riak AQ LoadAve



So, what is “at scale” ?

trends and behavior, not events

flexibility in deployment through configuration management

good operational resources

cache layers

structural homogeneity

evenly balanced

repeatable operational procedures

scale is step by step

like a musical tonality ::: precisely defined ::: expandable

