PID loops and the art of keeping systems stable

@colmmacc

Colm MacCárthaigh

2019-06-24

Control Theory: Where the fruit is hanging so low IT IS TOUCHING THE GROUND

META

aws

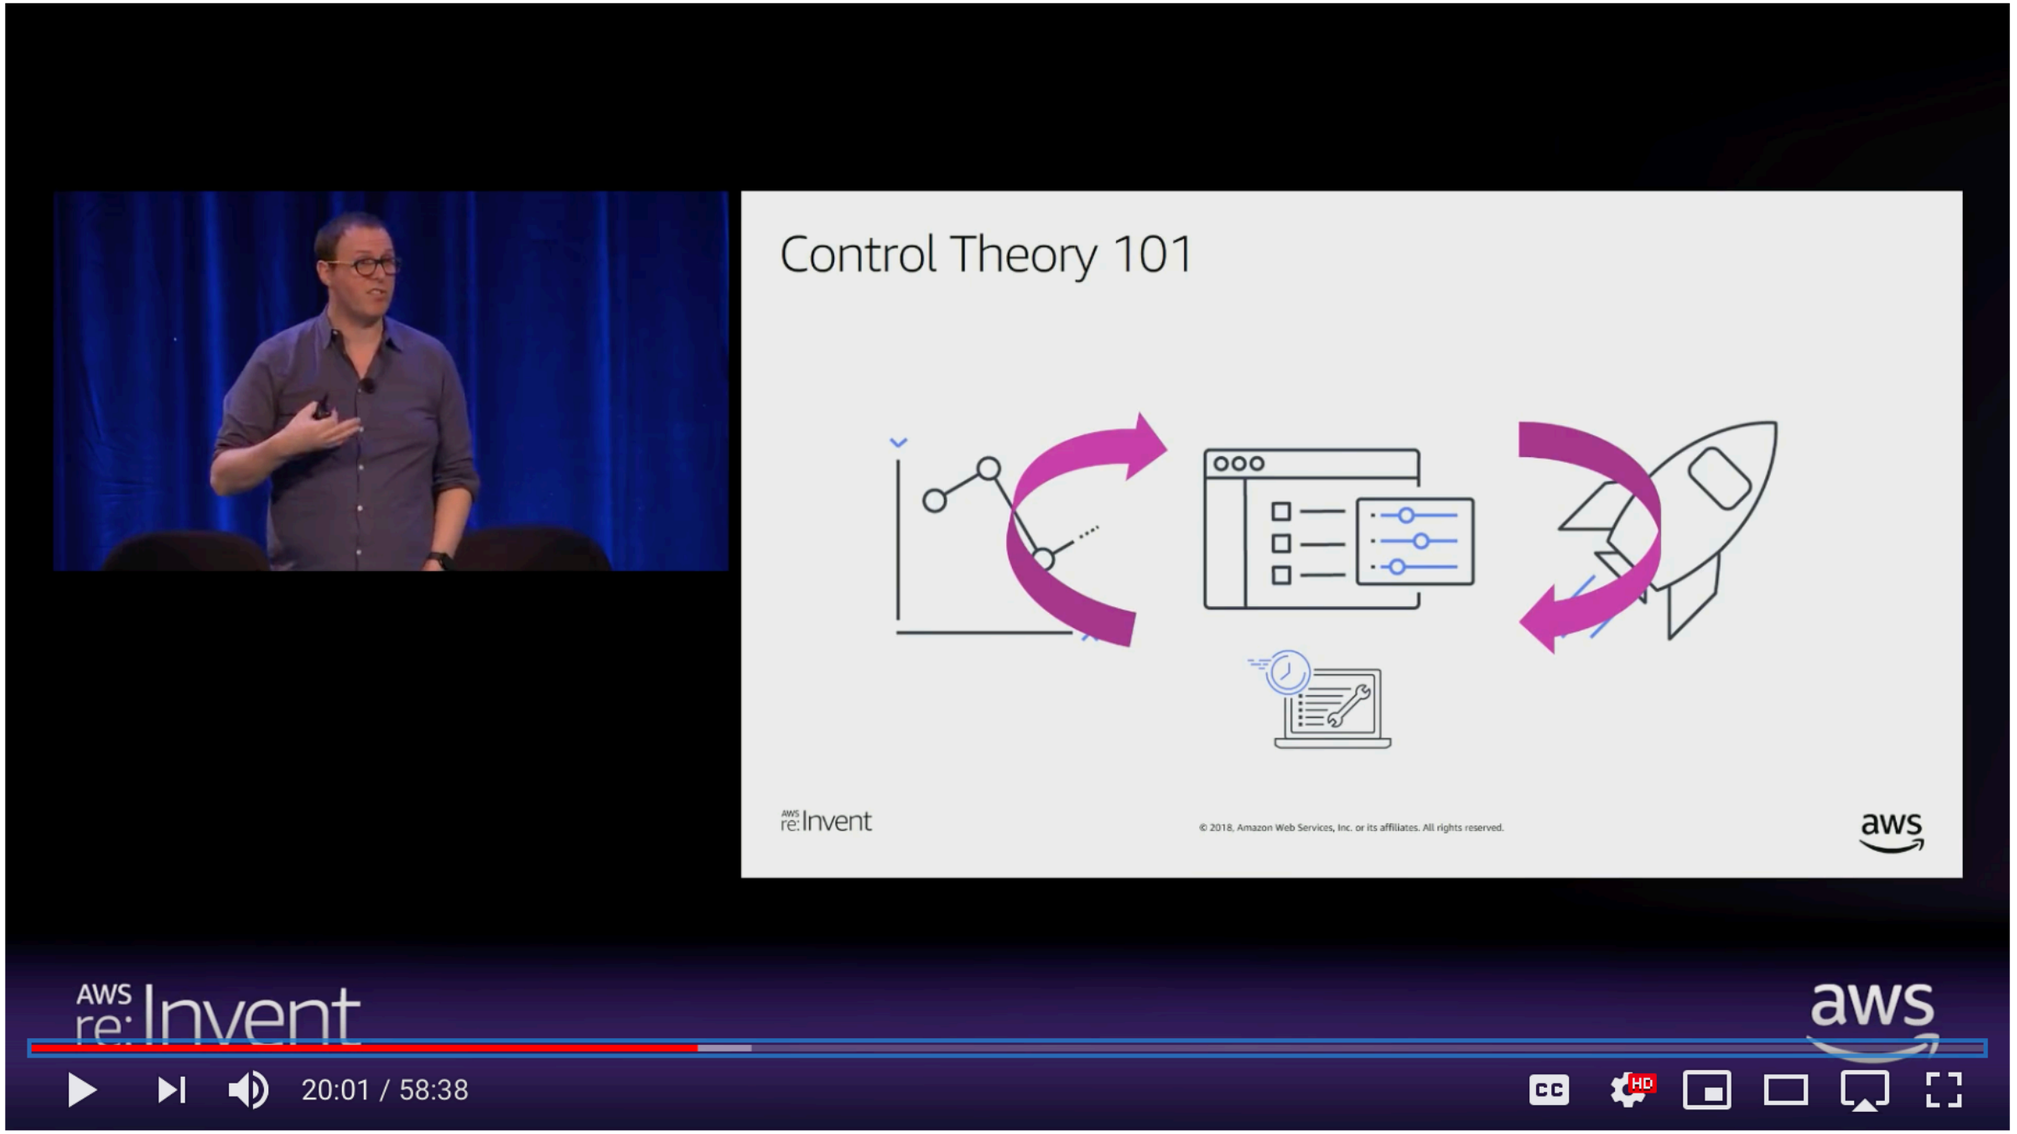

AWS re:Invent 2018: Close Loops & Opening Minds: How to Take Control of Systems, Big & Small ARC337

Hello 👋

Apologies in advance for the wall of text!

I’m hosting the “Modern CS in the real world” track at a conference called QCon in NYC this June. I was wondering if you might be interested in giving a talk about PID loops and their application in building distributed systems. I loved the references to the idea in some of your previous talks - but it’d be awesome to see some of those ideas be discussed more in depth.

Let me know.

Mar 7

I'd love to and I have a lot more depth that wasn't in the re:Invent talk!

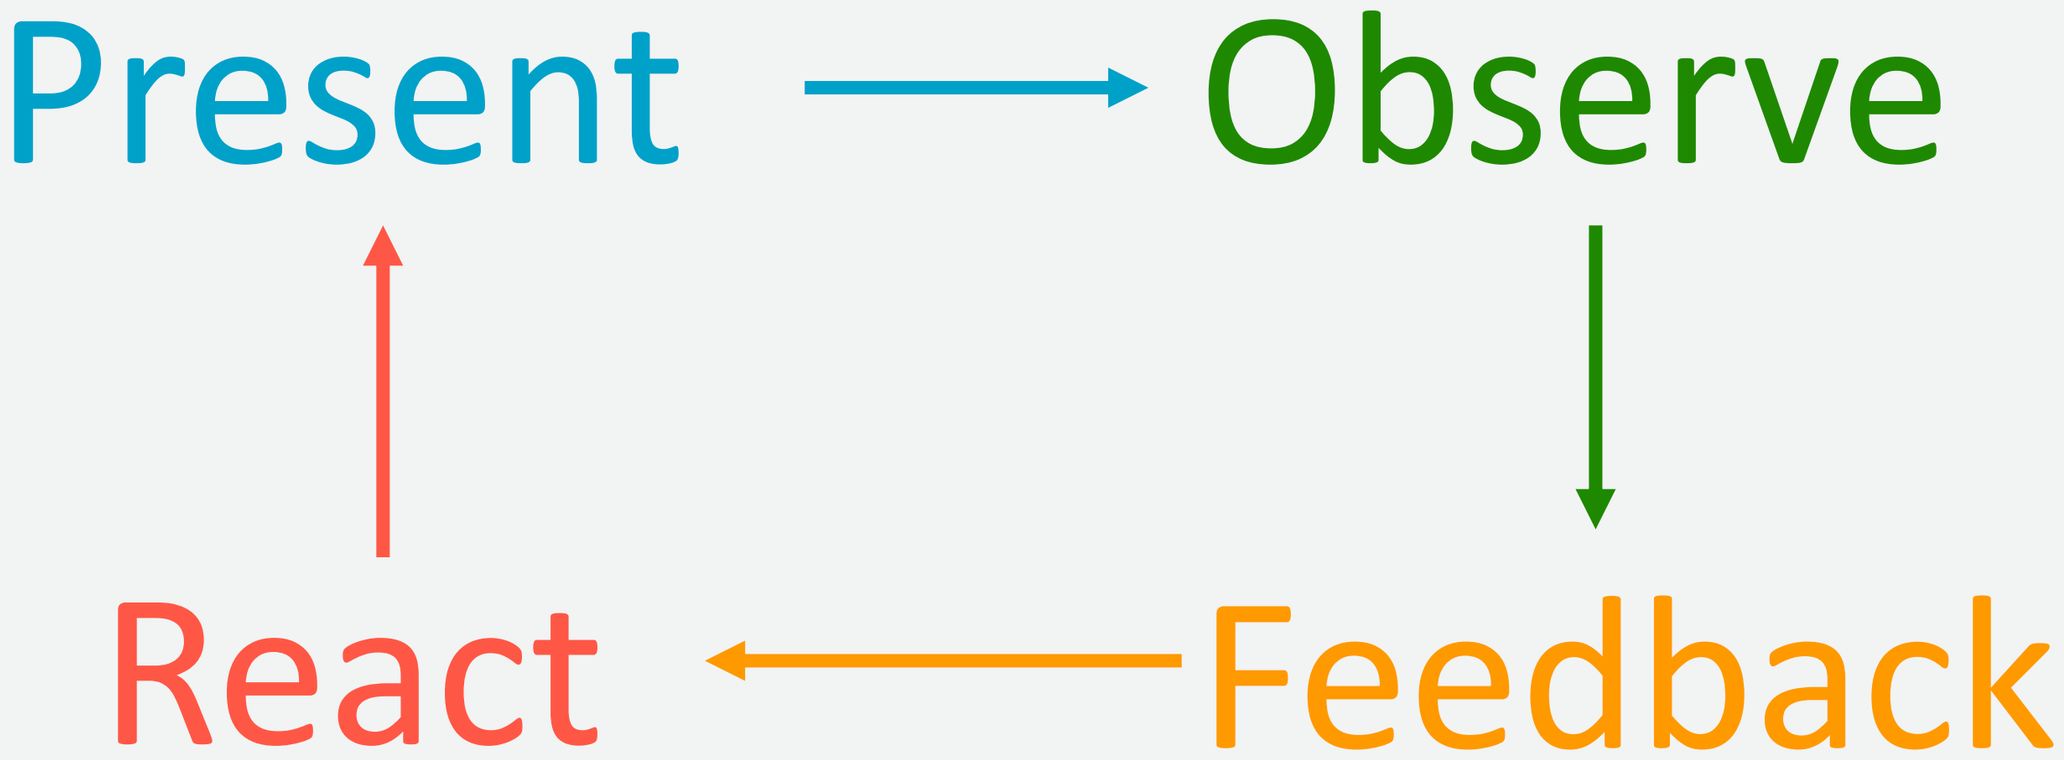

aws

Control Theory

aws

Prior Art

Control Theory in Container Fleet Management

LIKE DISCUSS BOOKMARKS

1X 1.5X 2X

MP3 SLIDES ANDROID APP 37:02

Summary

Vallery Lancey covers basic principles of observing systems, controller design, and PID controllers. In particular, she dives into container scaling controllers, using both first principles and proven designs from Kubernetes and Mesos.

```
while True {
    currentState = getCurrentState()
    desiredState = getDesiredState()
    makeConform(currentState, desiredState)
}
```

Setpoint

Process Variable

@vllry

aws

Prior Art

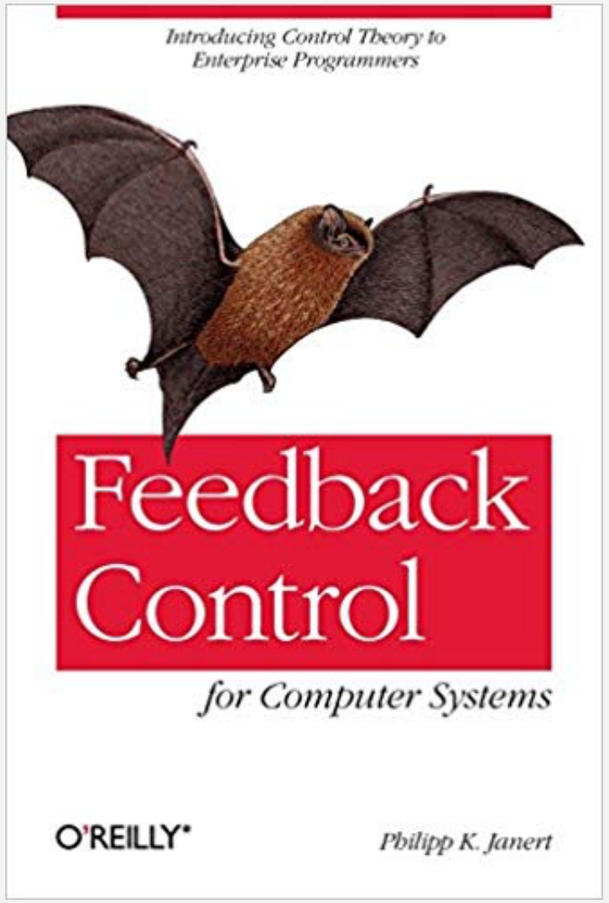

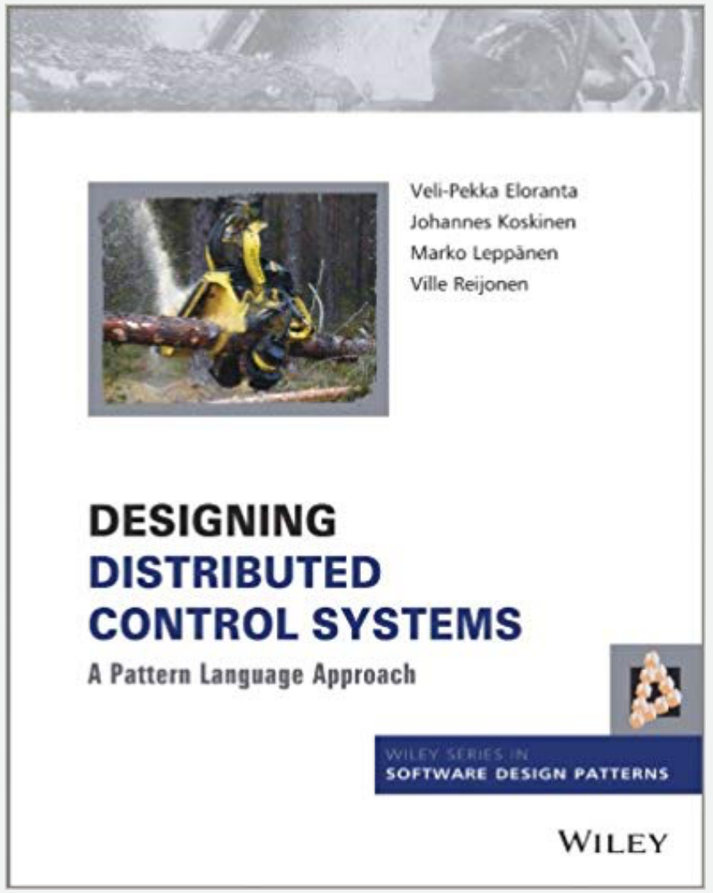

aws

Control Theory and PID loops

Comes up in the context of …

Autoscaling and placement: Instances, Storage, Network, etc.

Fairness algorithms: TCP, Queues, Throttling

Systems stability

The Furnace

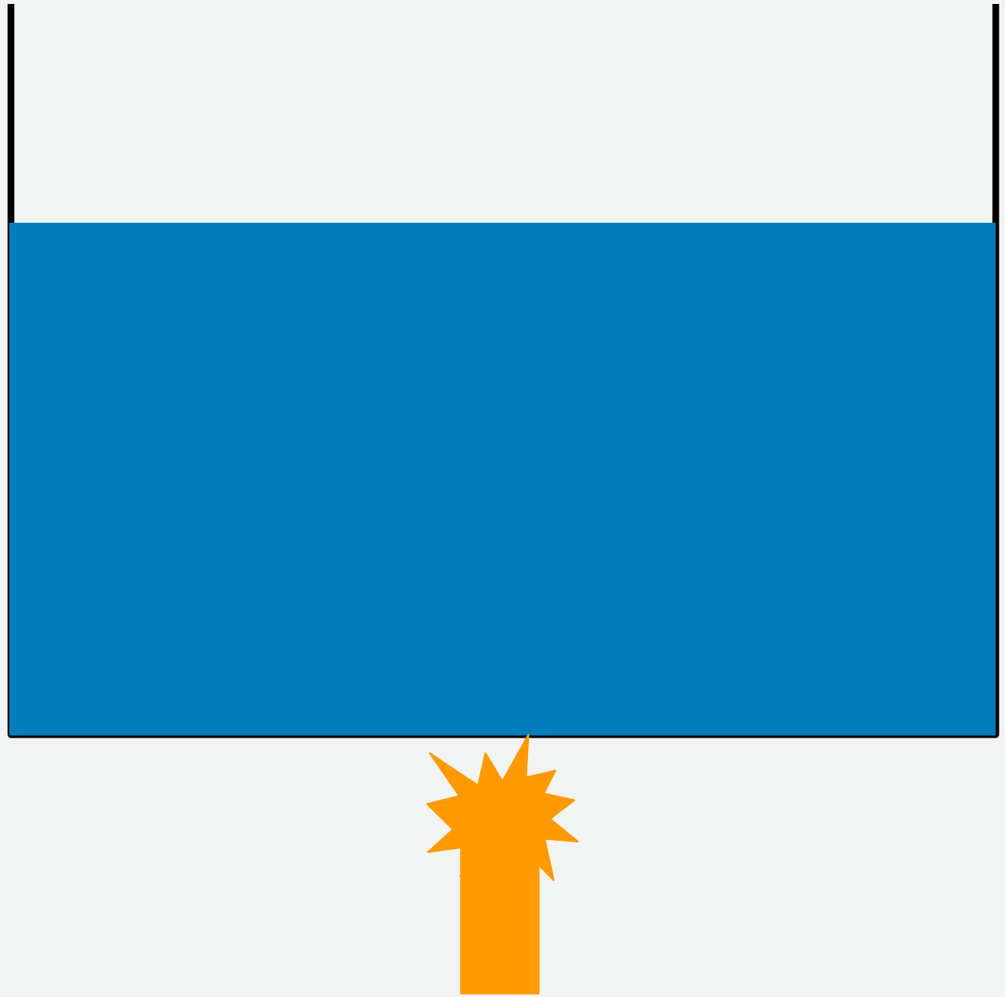

The Furnace

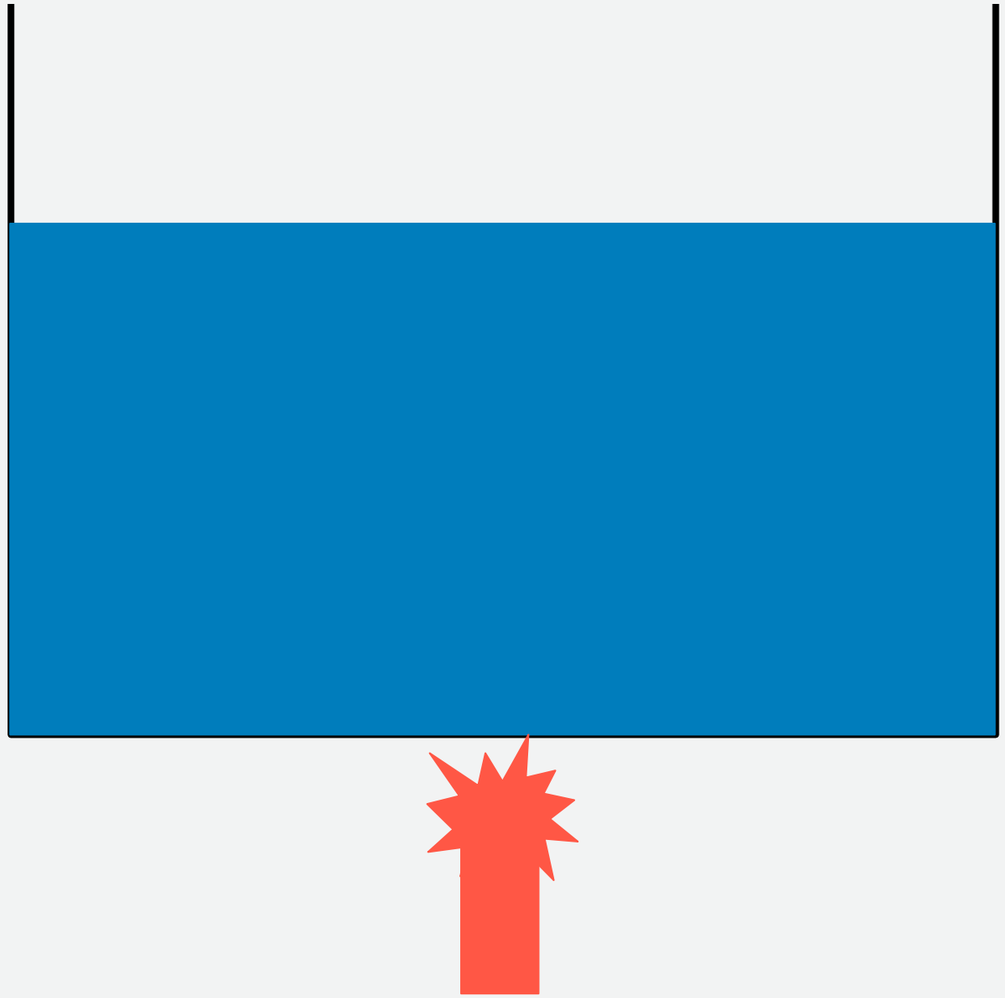

The Furnace

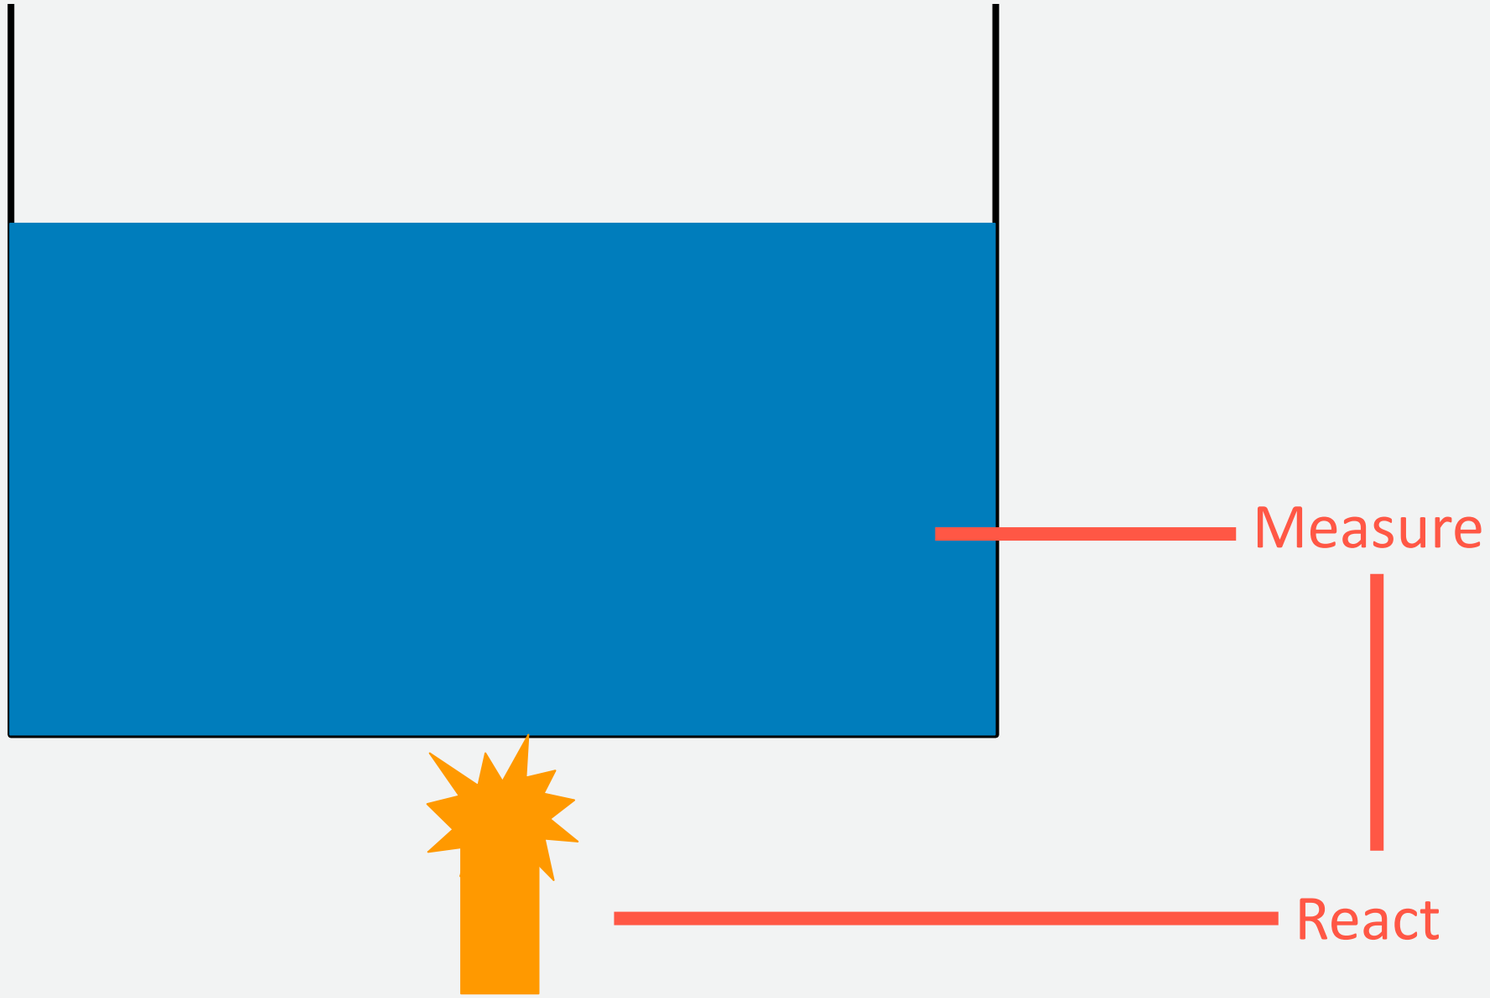

The Furnace

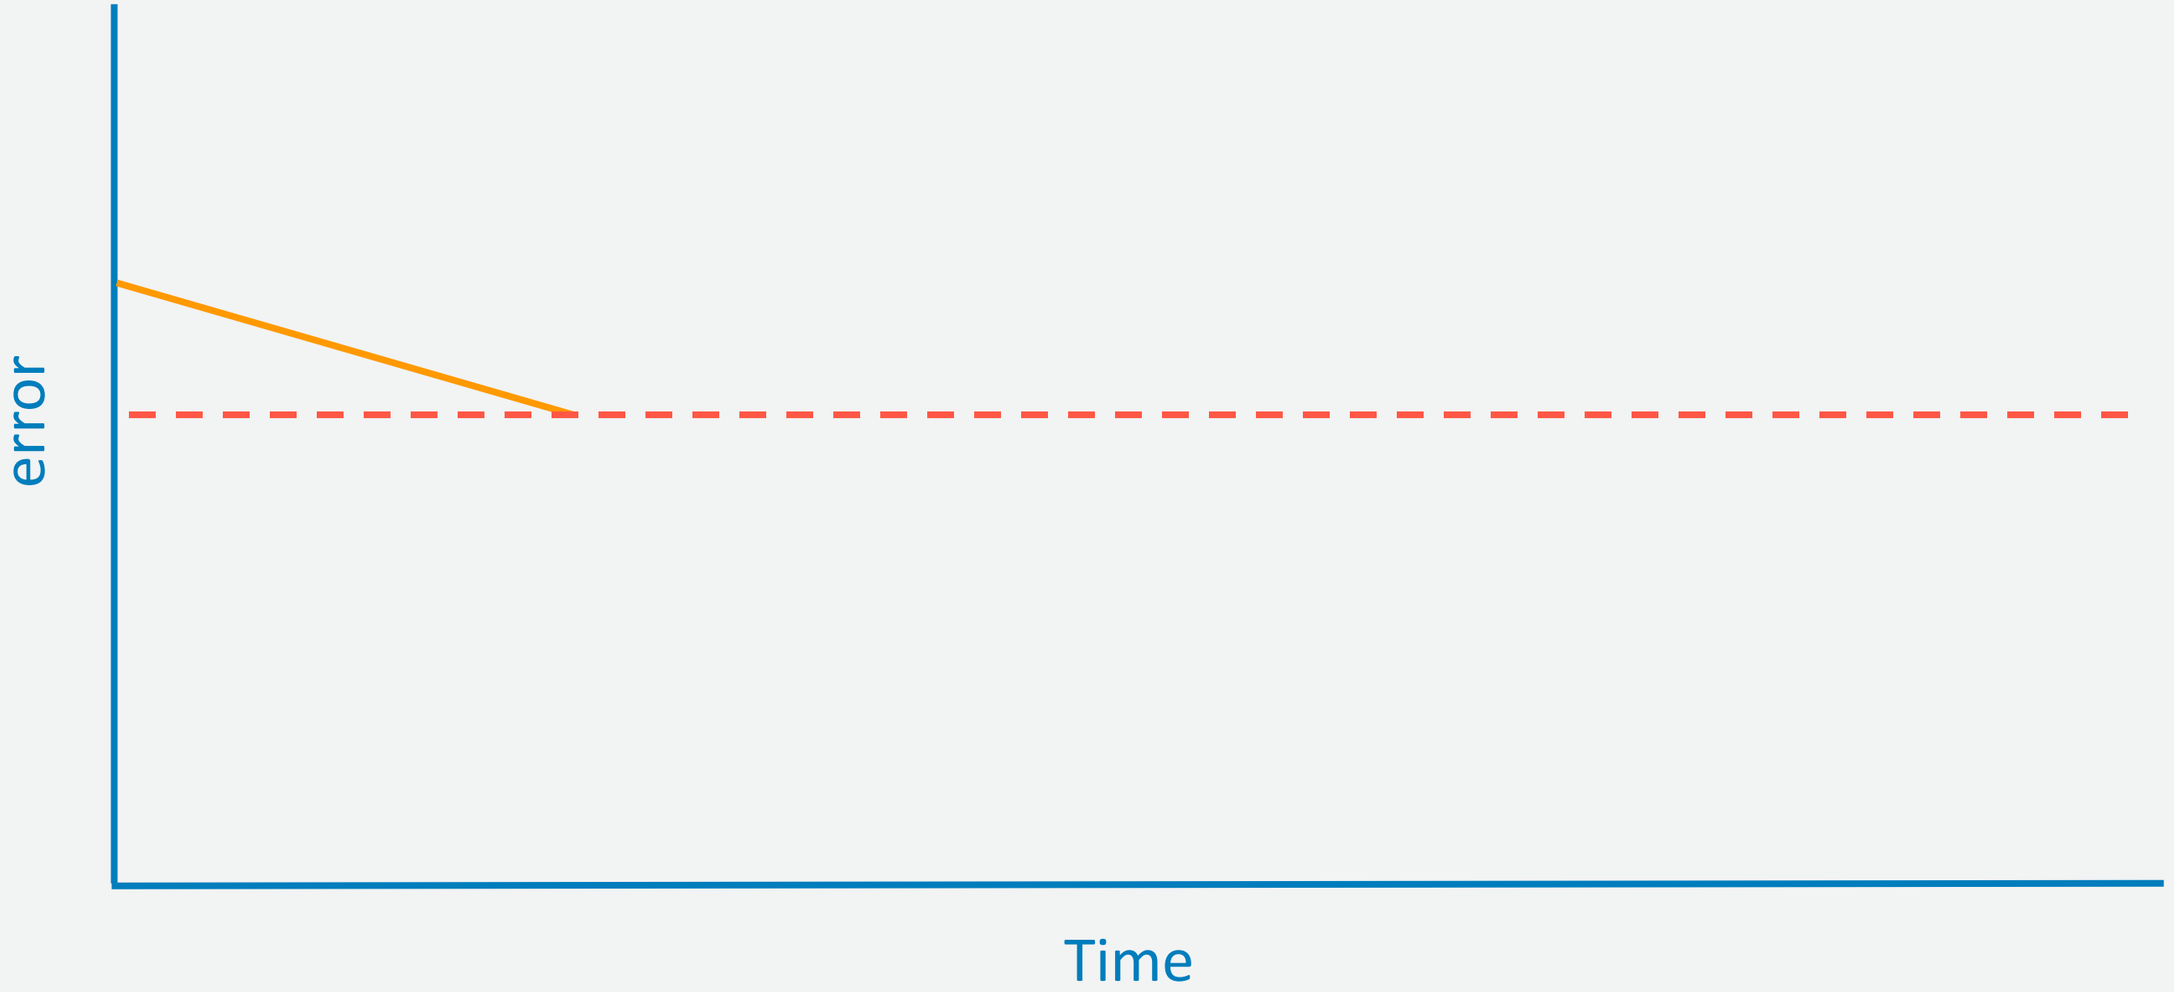

aws

The Furnace

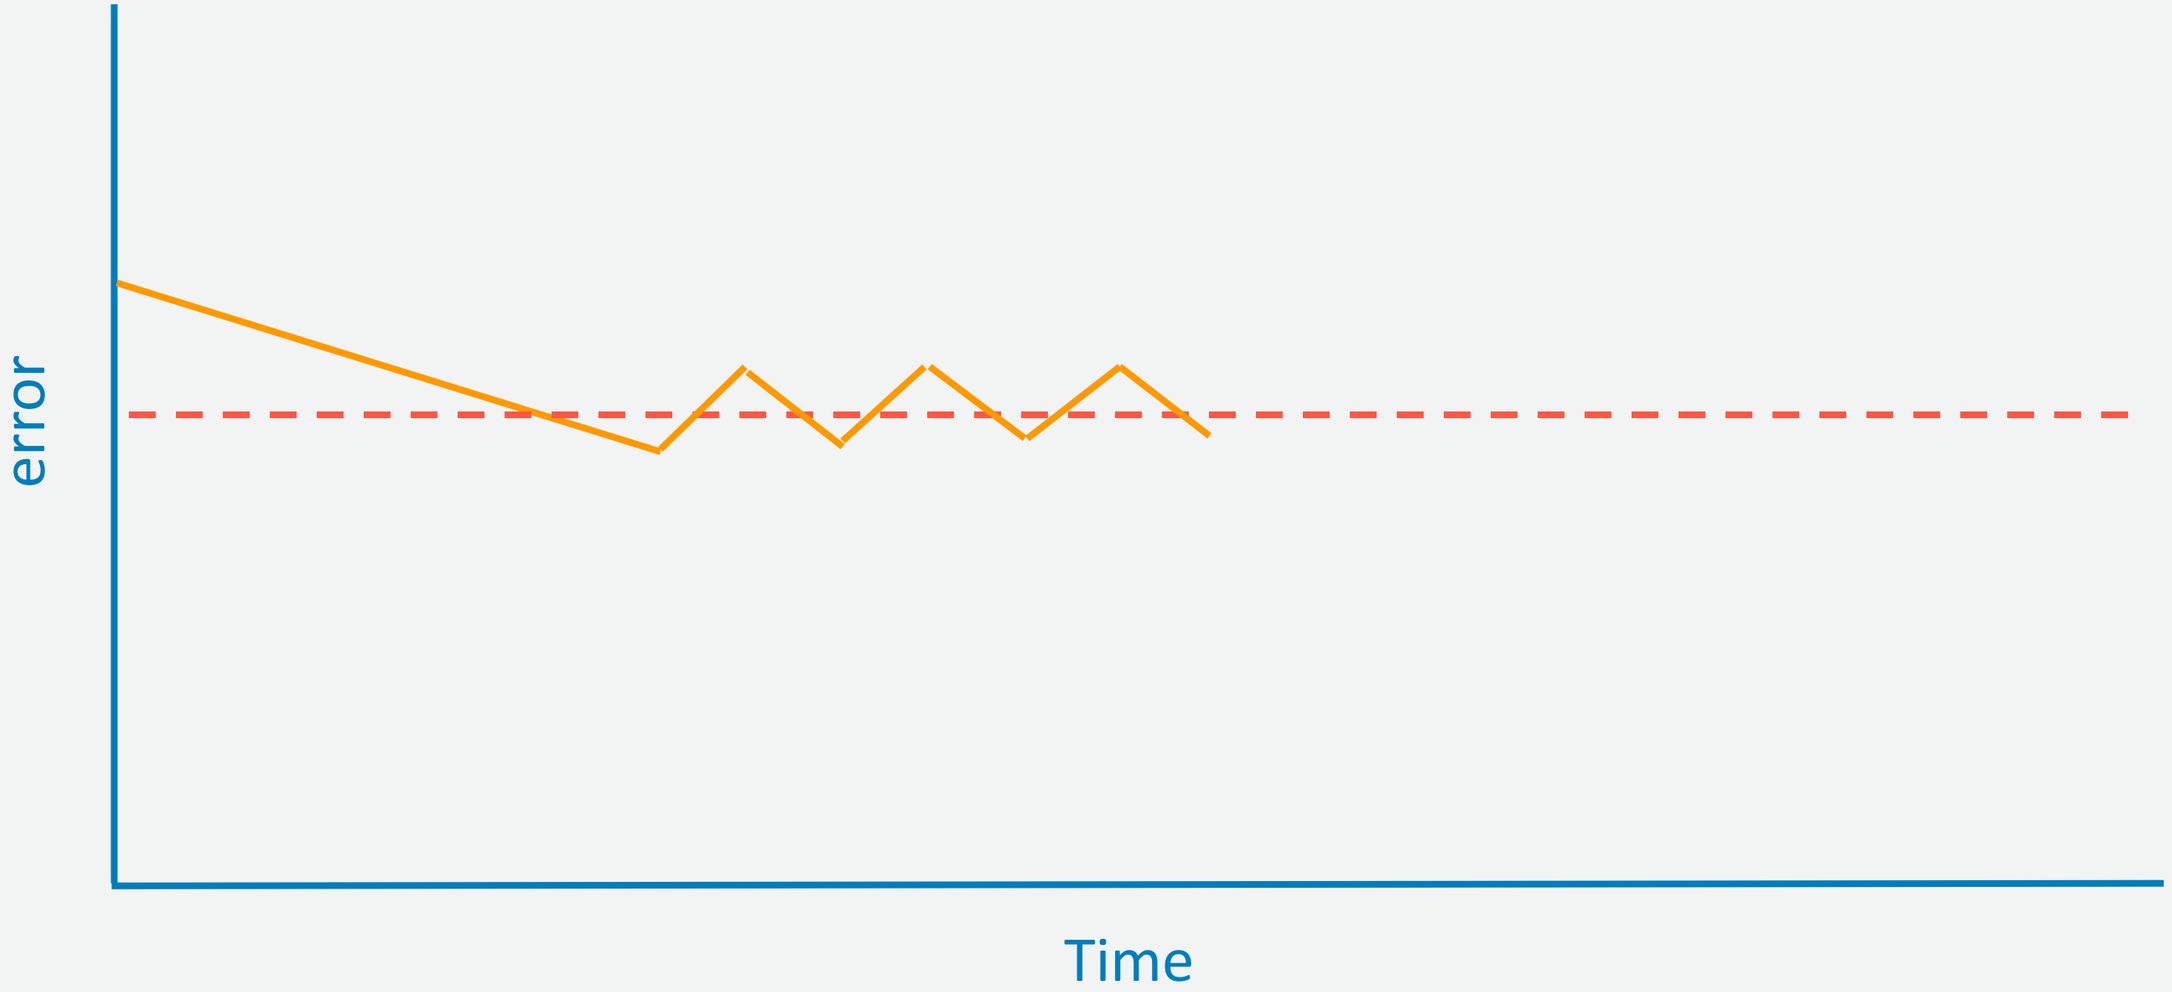

aws

The Furnace

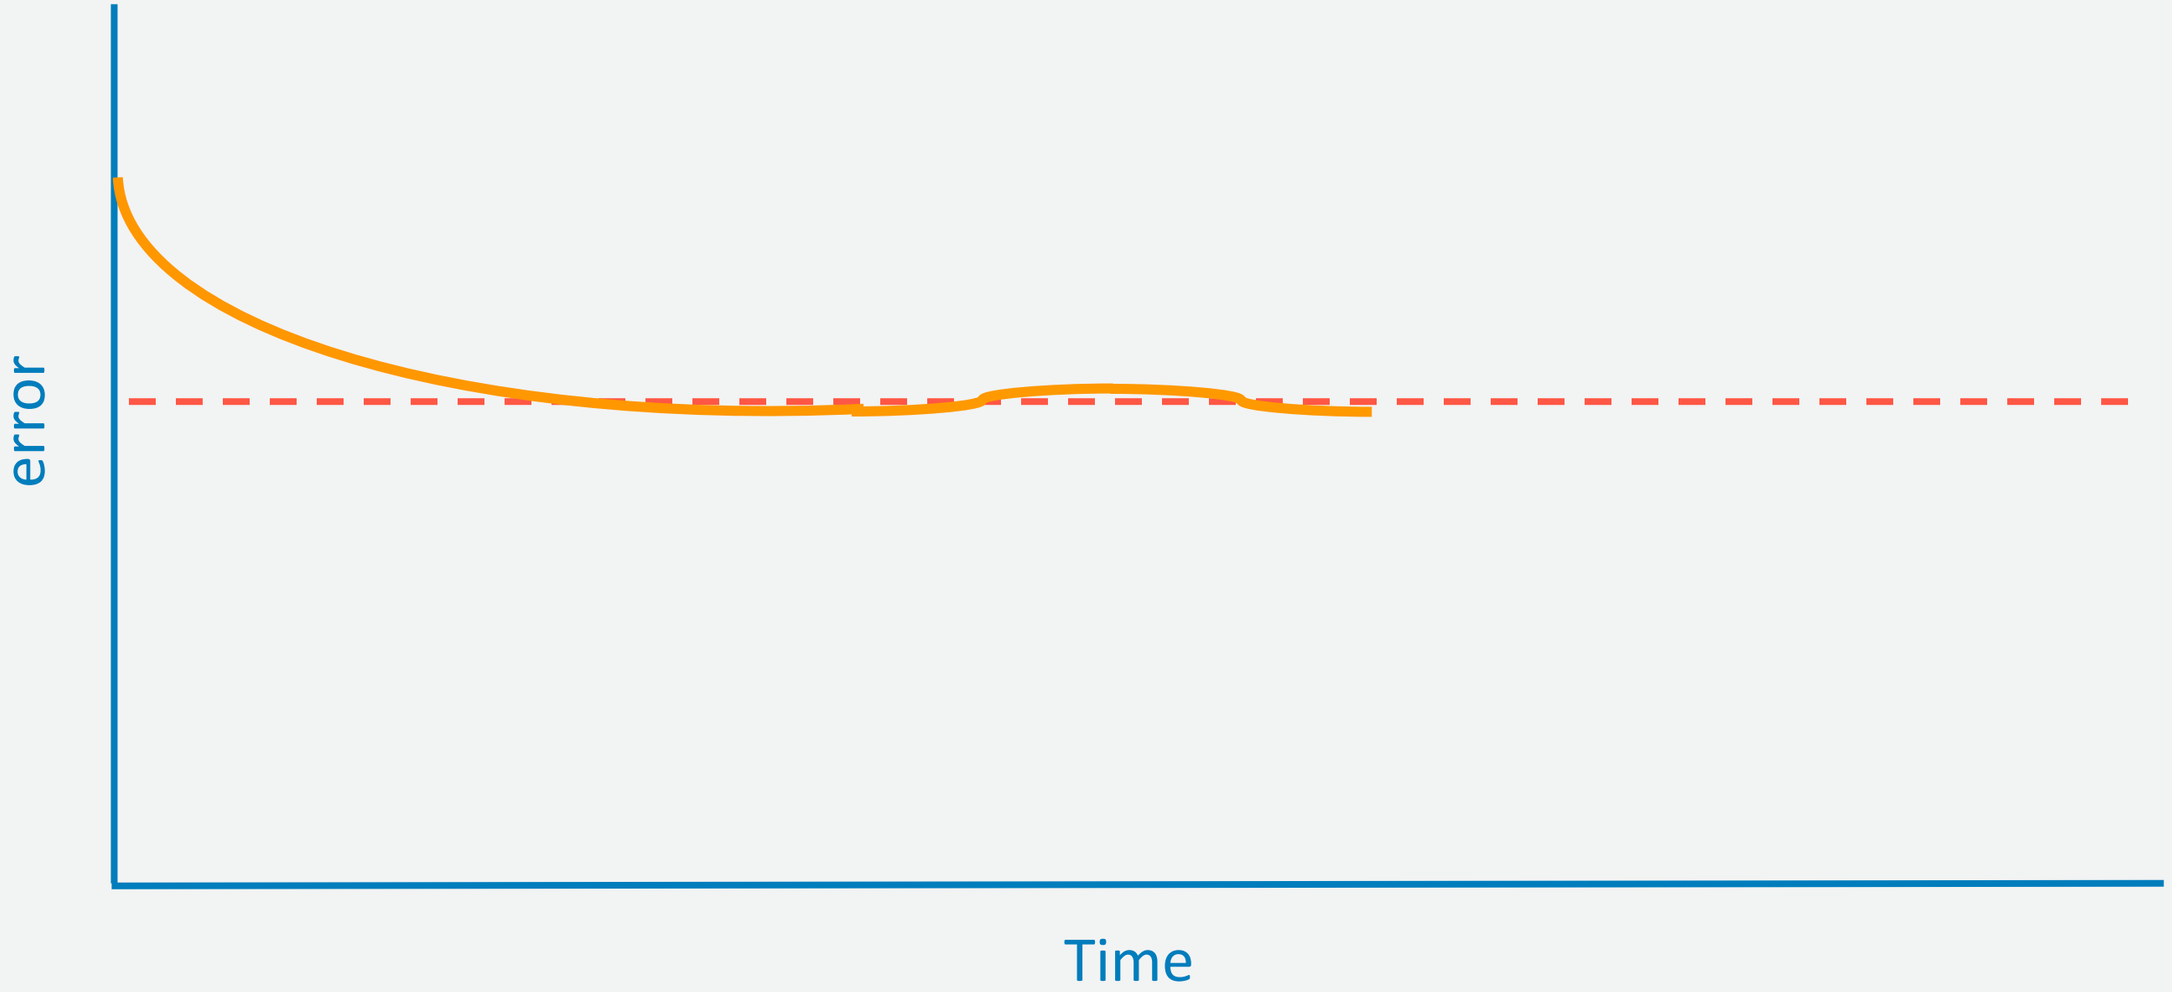

aws

Autoscaling

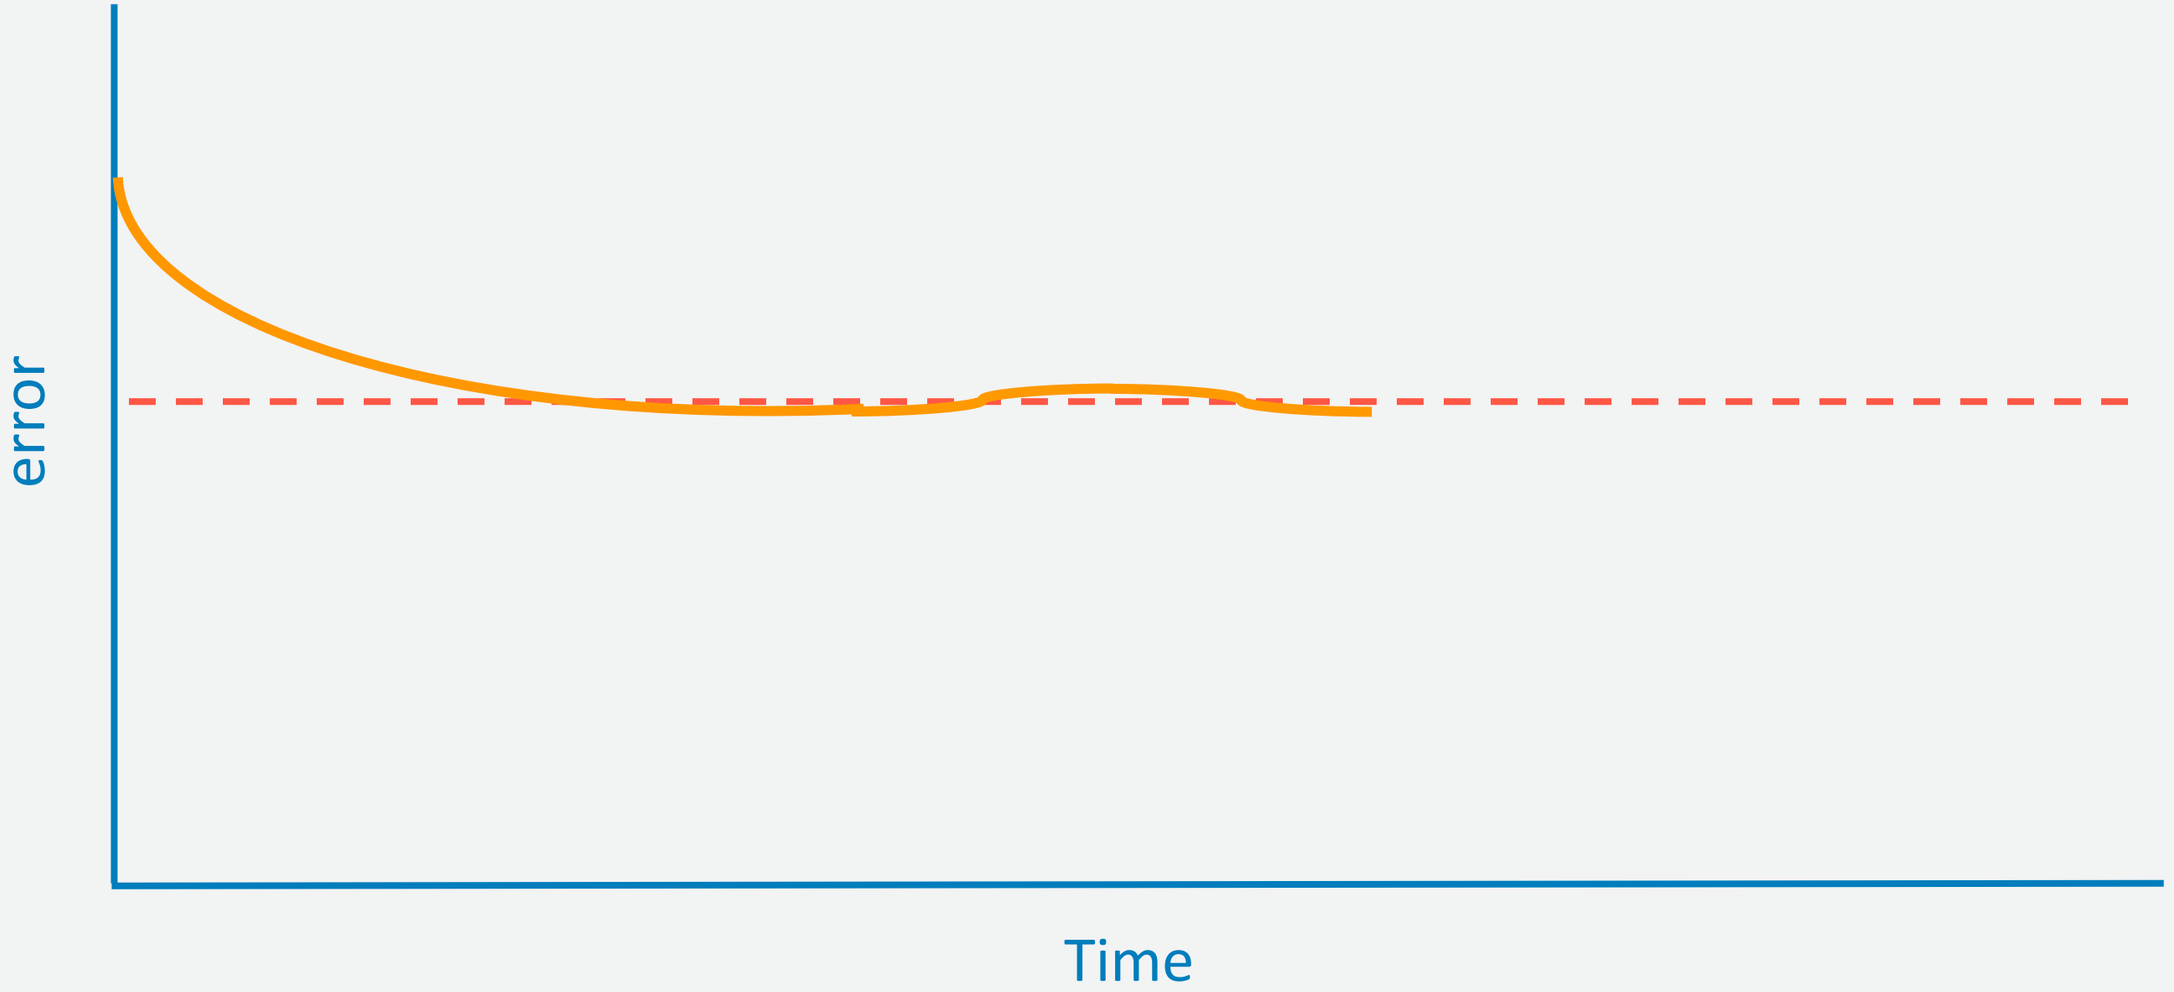

aws

Autoscaling

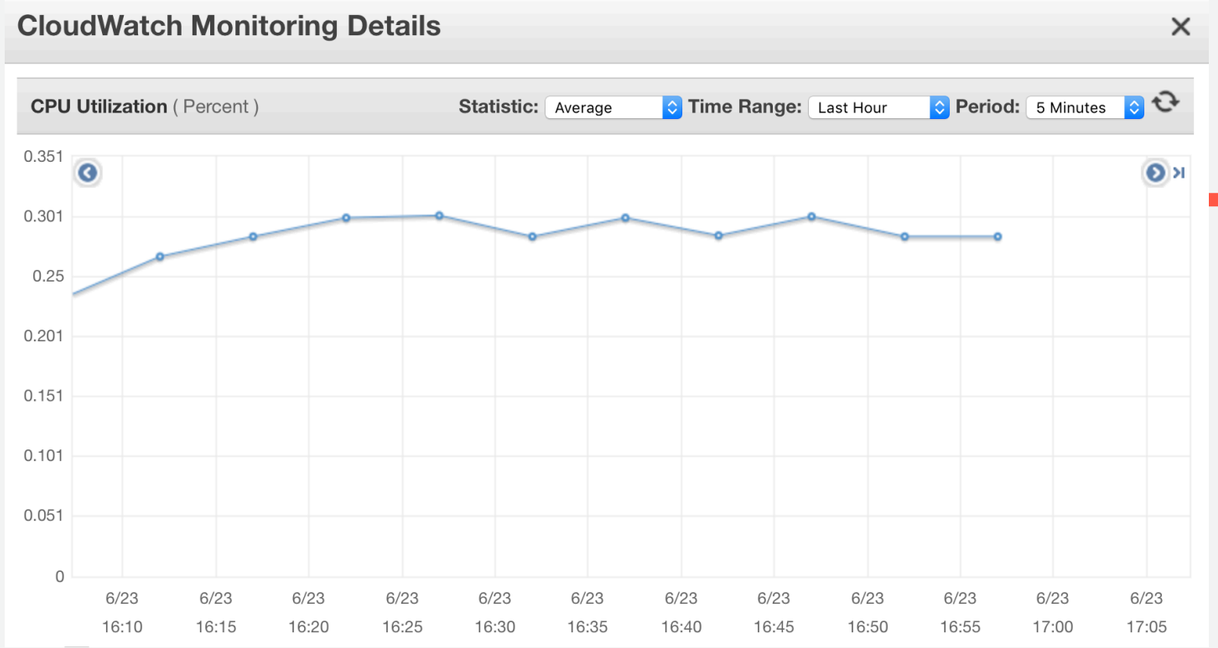

Measure

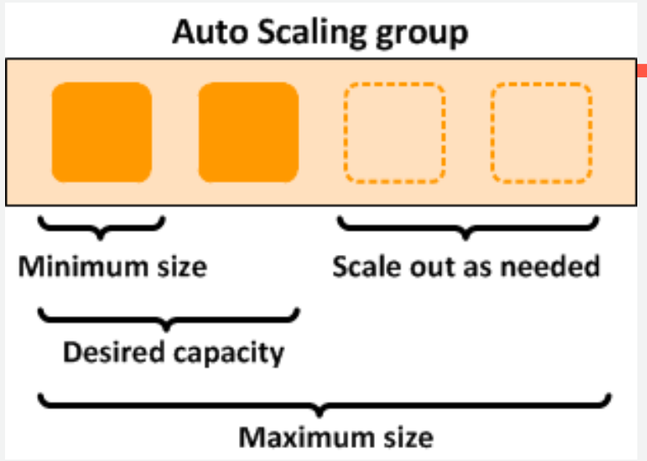

React

aws

Autoscaling : forecasting and fancy integrals!

Any signal can be processed with Fourier Analysis to find underlying constituent frequencies

Real-world operational systems often have strong daily, weekly, annual cycles, etc.

Holt-Winters Forecasting can simulate these cycles into the future

Machine Learning can do even better!

aws

Autoscaling : forecasting and fancy integrals!

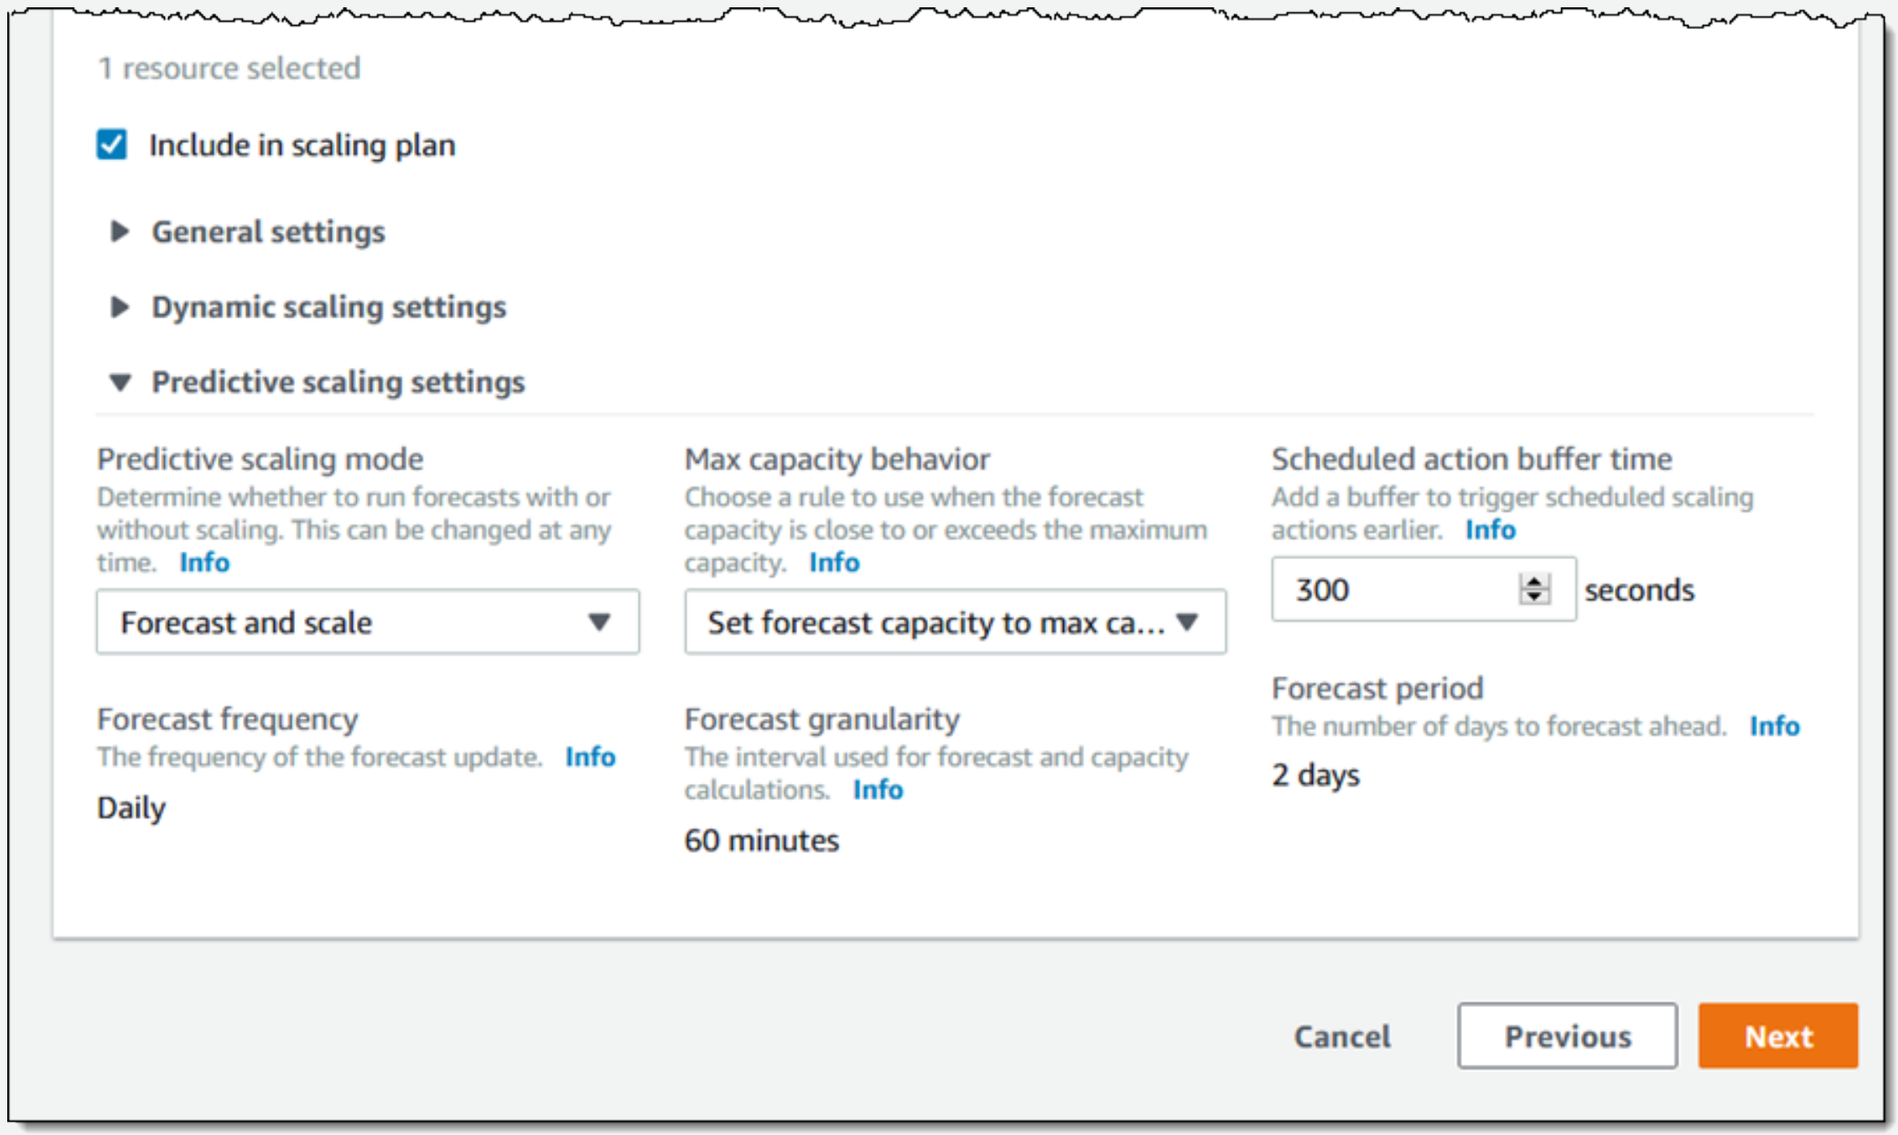

Placement and fairness

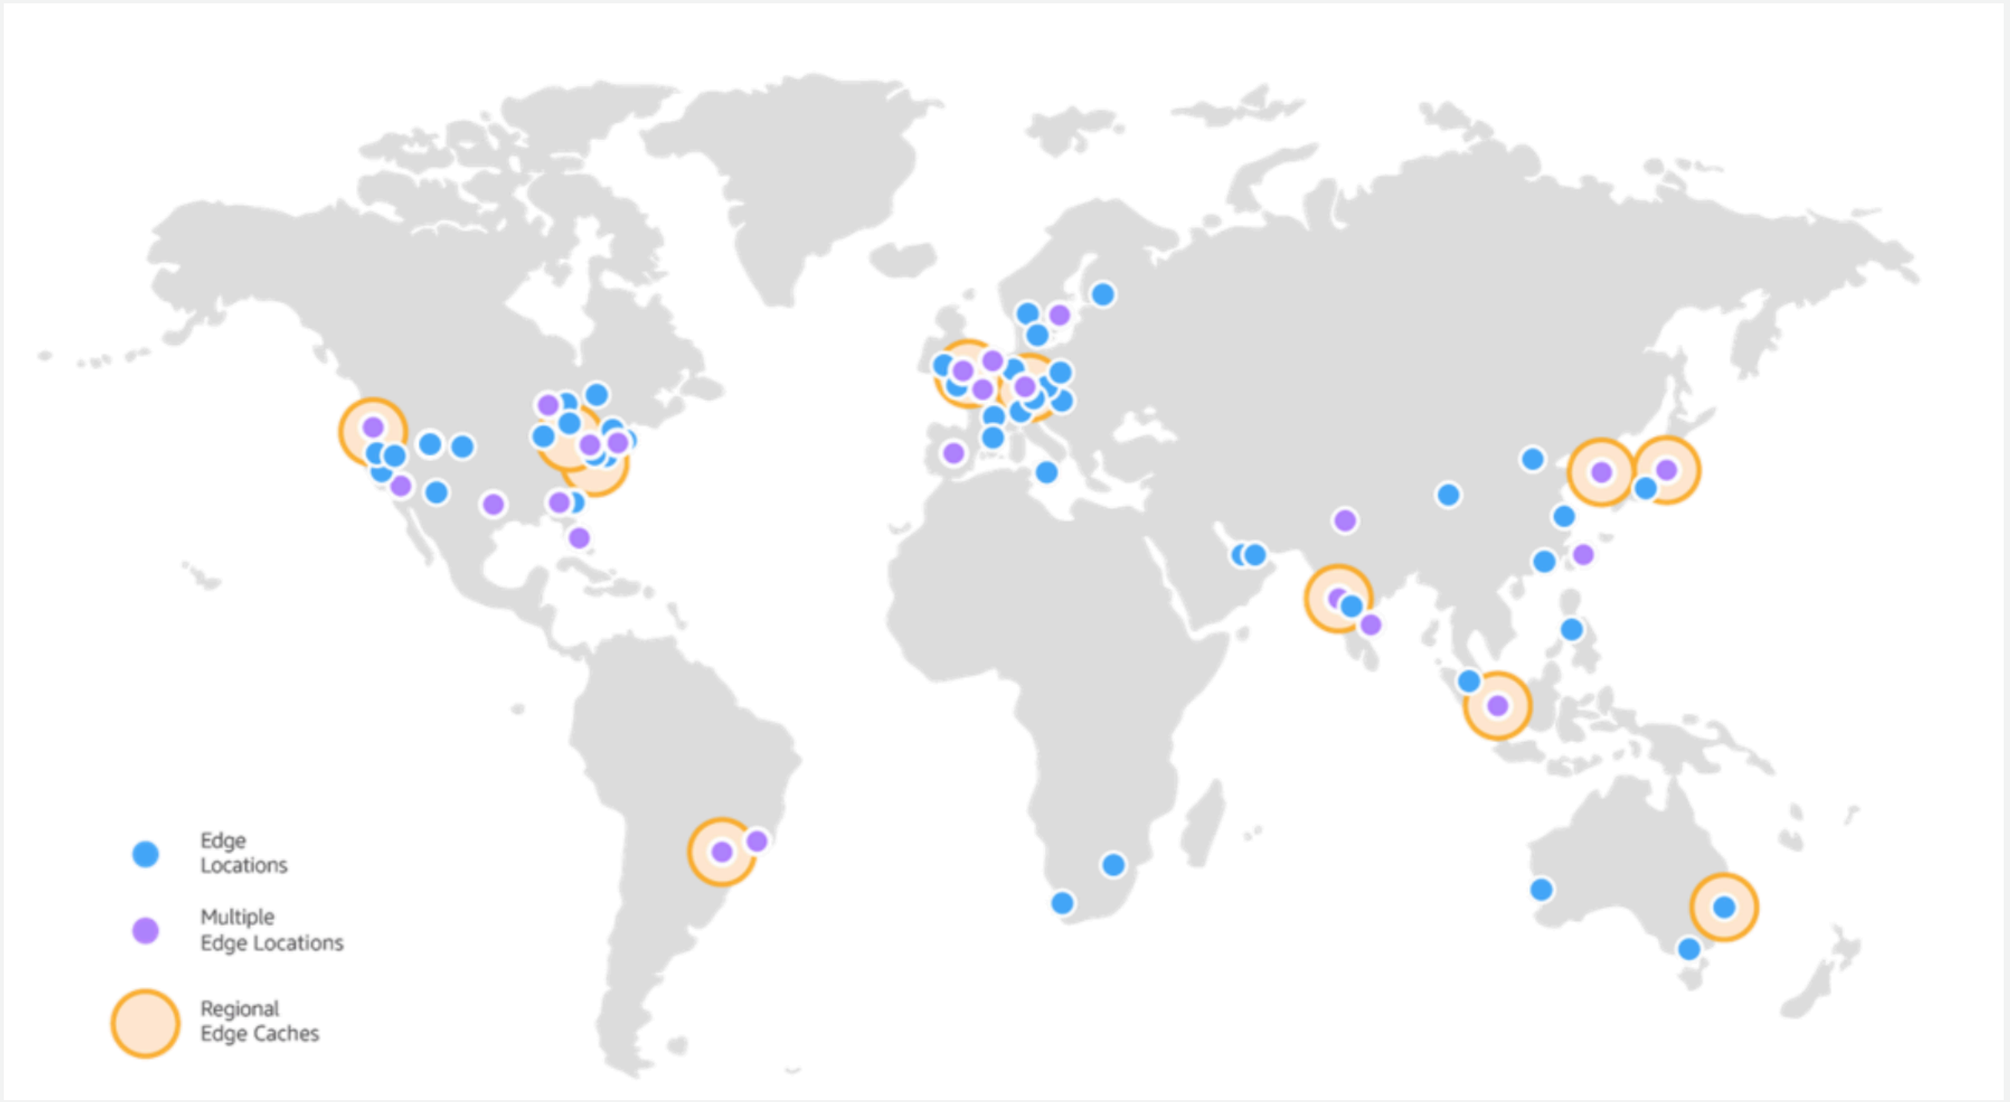

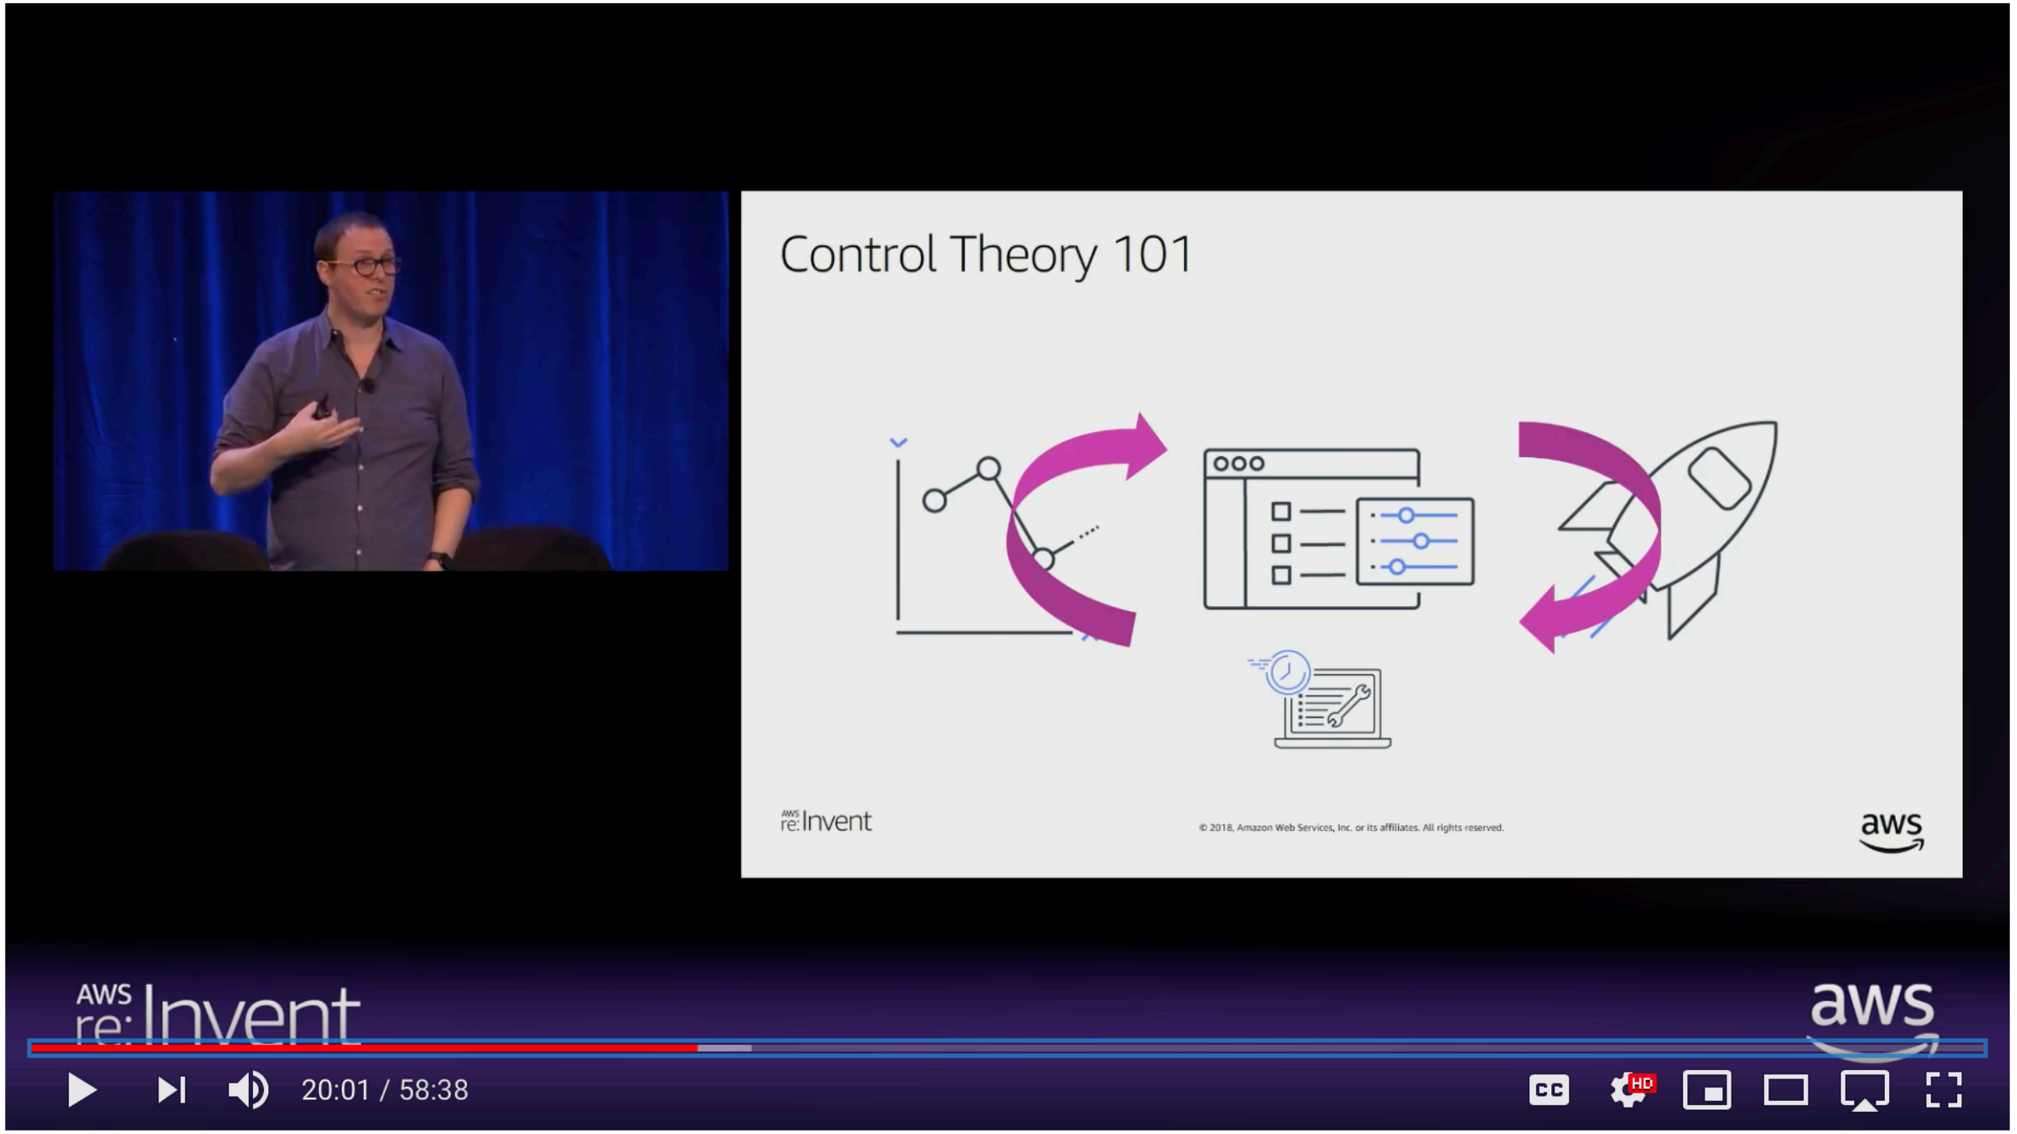

AWS re:Invent 2018: Close Loops & Opening Minds: How to Take Control of Systems, Big & Small ARC337

1

X-Ray Vision: Open Loops

aws

X-Ray Vision: Open Loops

```
# launch 10 Instances
for (i = 0; i < 10; i++)
    instance[i] = ec2_launch_instance()

# wait a minute
sleep(60);

# Register the instances
for (i = 0; i < 10; i++)
    register_instance(instance[i]);
```


aws

X-Ray Vision: Open Loops

- A surprising number of real-world systems are Open Loops
- Potential reasons:
    - Organic out-growth from scripts
    - Imperative programming "*Do this, then do this*" is very natural
    - Infrastructure is very very reliable these days
    - Infrequent actions

aws

X-Ray Vision: Open Loops

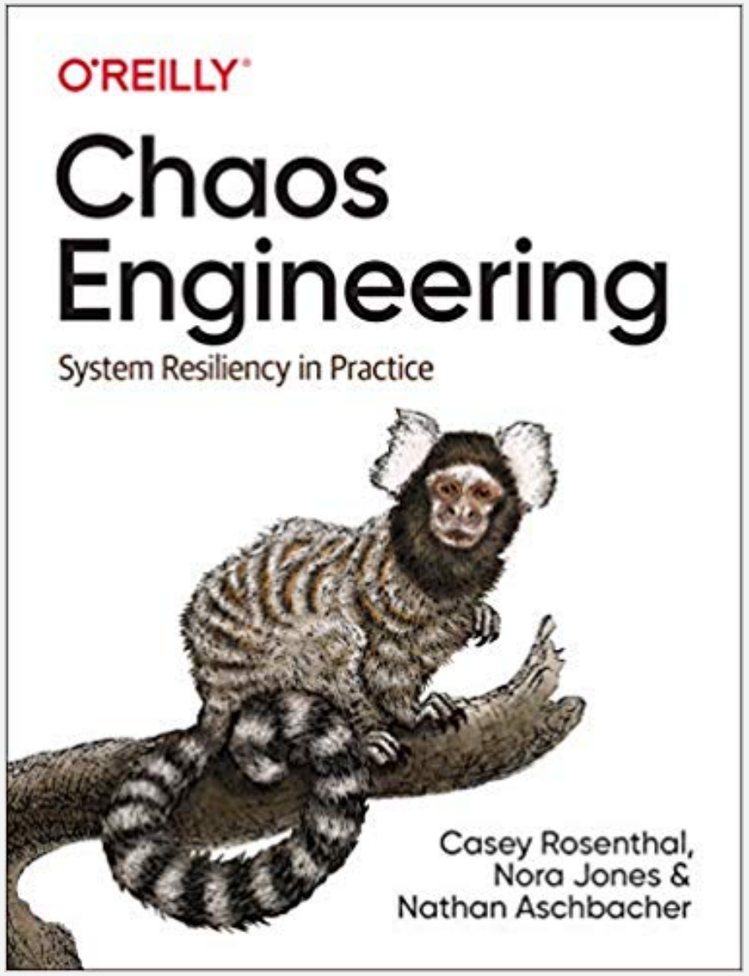

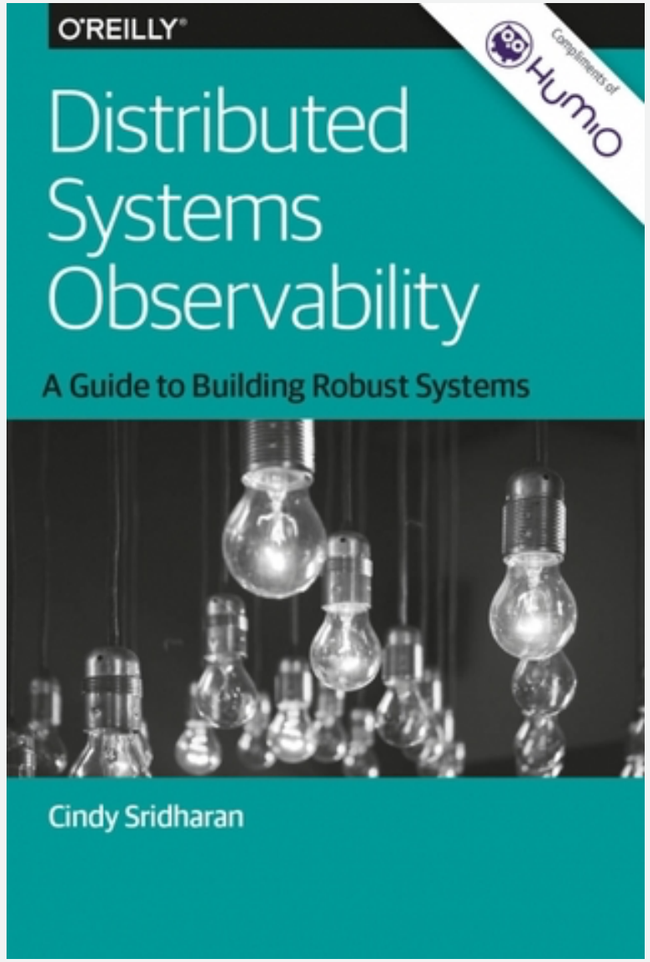

aws

X-Ray Vision: Open Loops

X-Ray Vision: Open Loops

- Closing loops:
    - Embrace "*Measure first. Then react.*"
    - Measure a lot of things. Check everything you can think to.
    - Avoid infrequent operations – make them more frequent where possible.

aws

X-Ray Vision: Open Loops

aws

X-Ray Vision: Open Loops

- Measure-first systems tend to be more naturally declarative
- In general, declarative are easily to formally verify
- TLA+ , F*, CoQ, SAW/Cryptol

aws

2

X-Ray Vision: Power Laws

aws

Power Laws

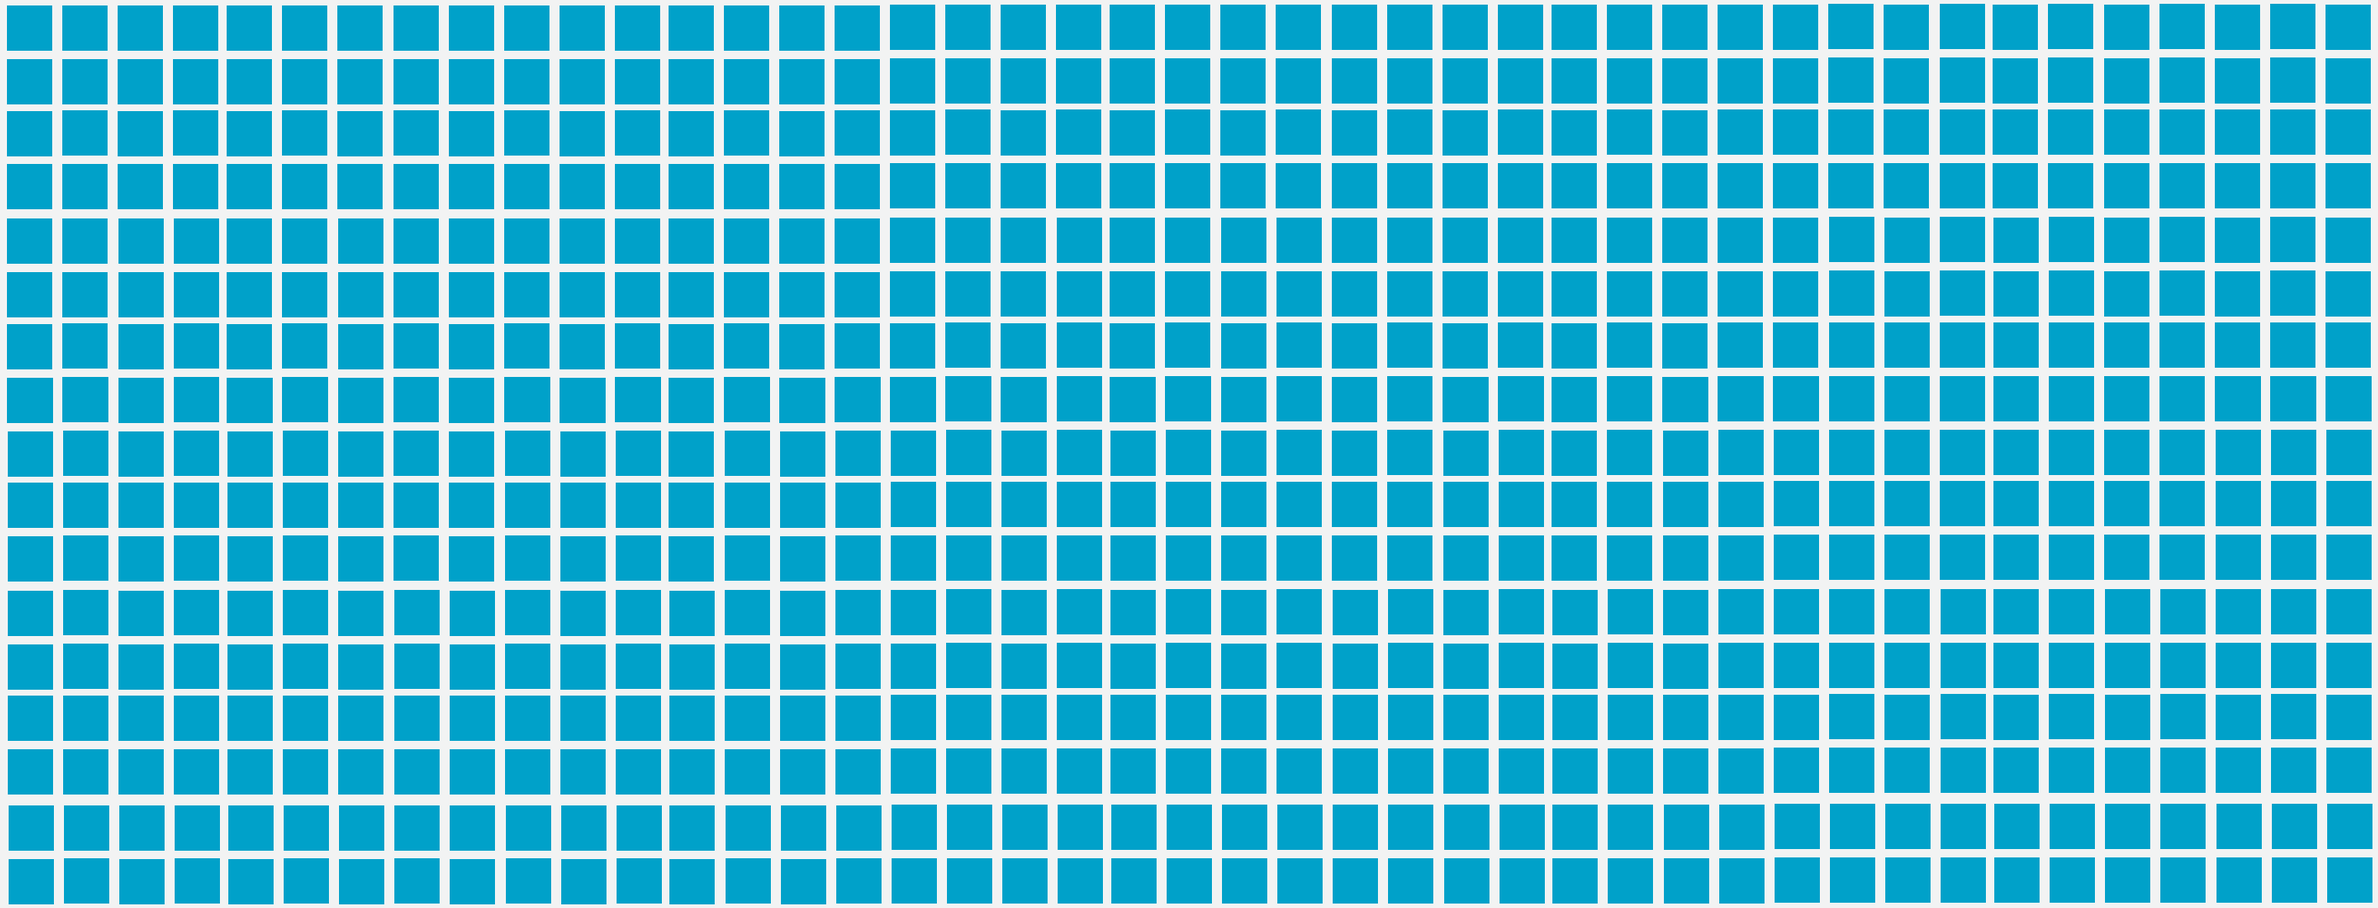

aws

Power Laws

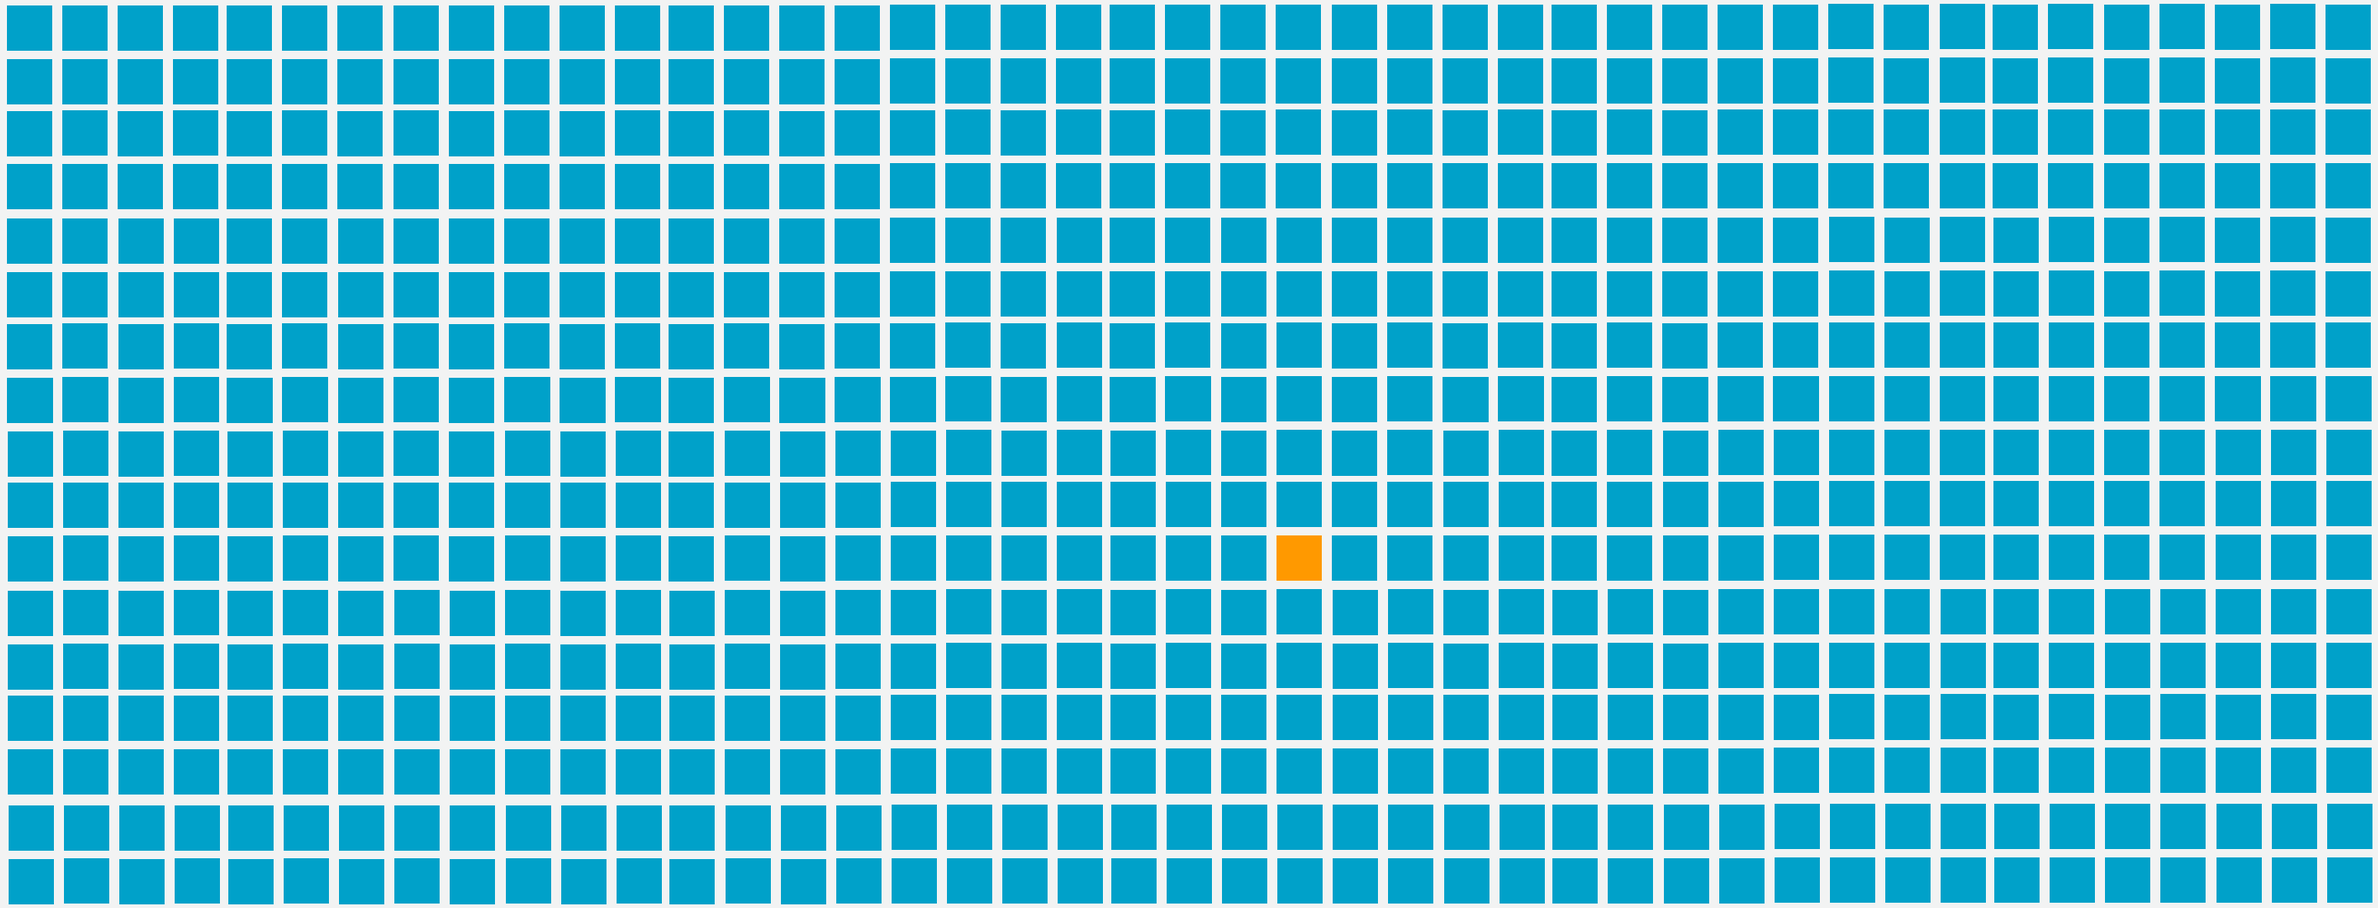

aws

Power Laws

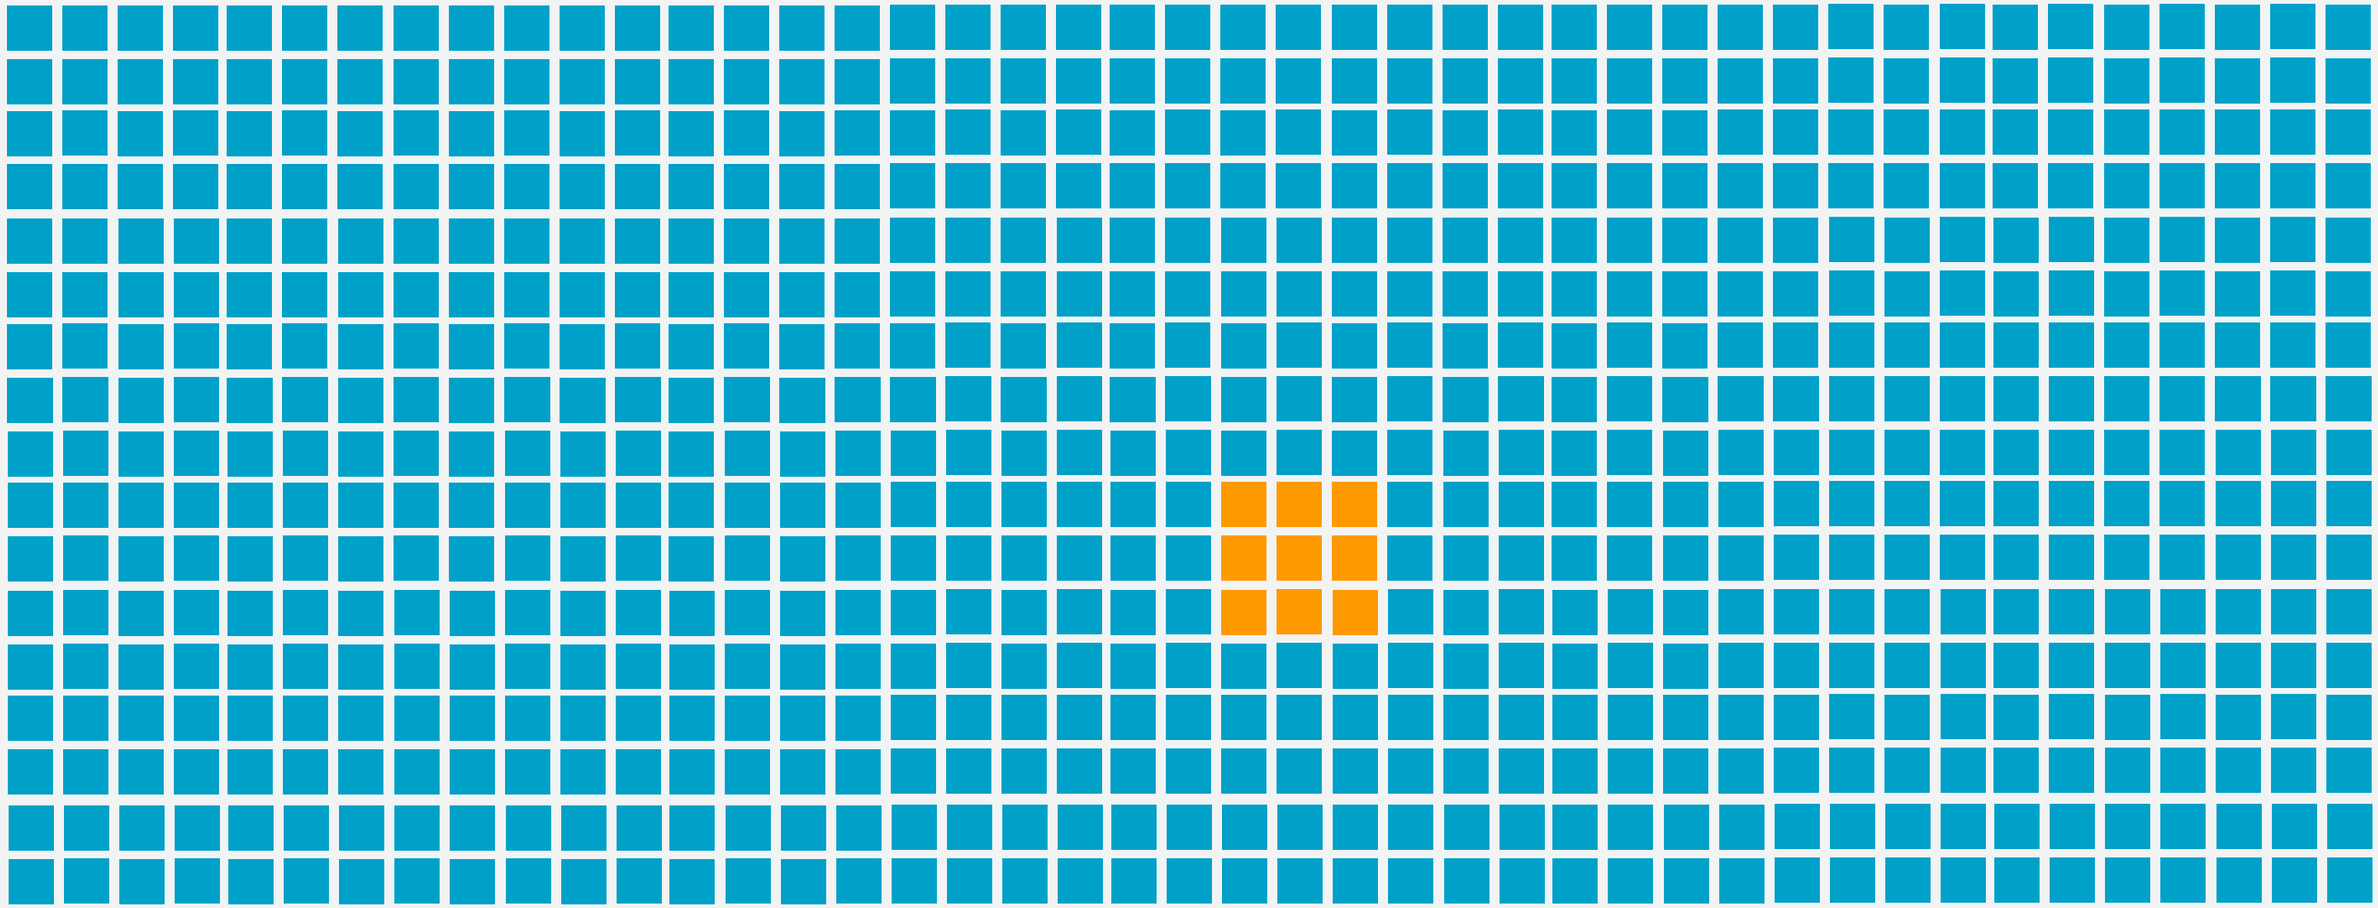

aws

Power Laws

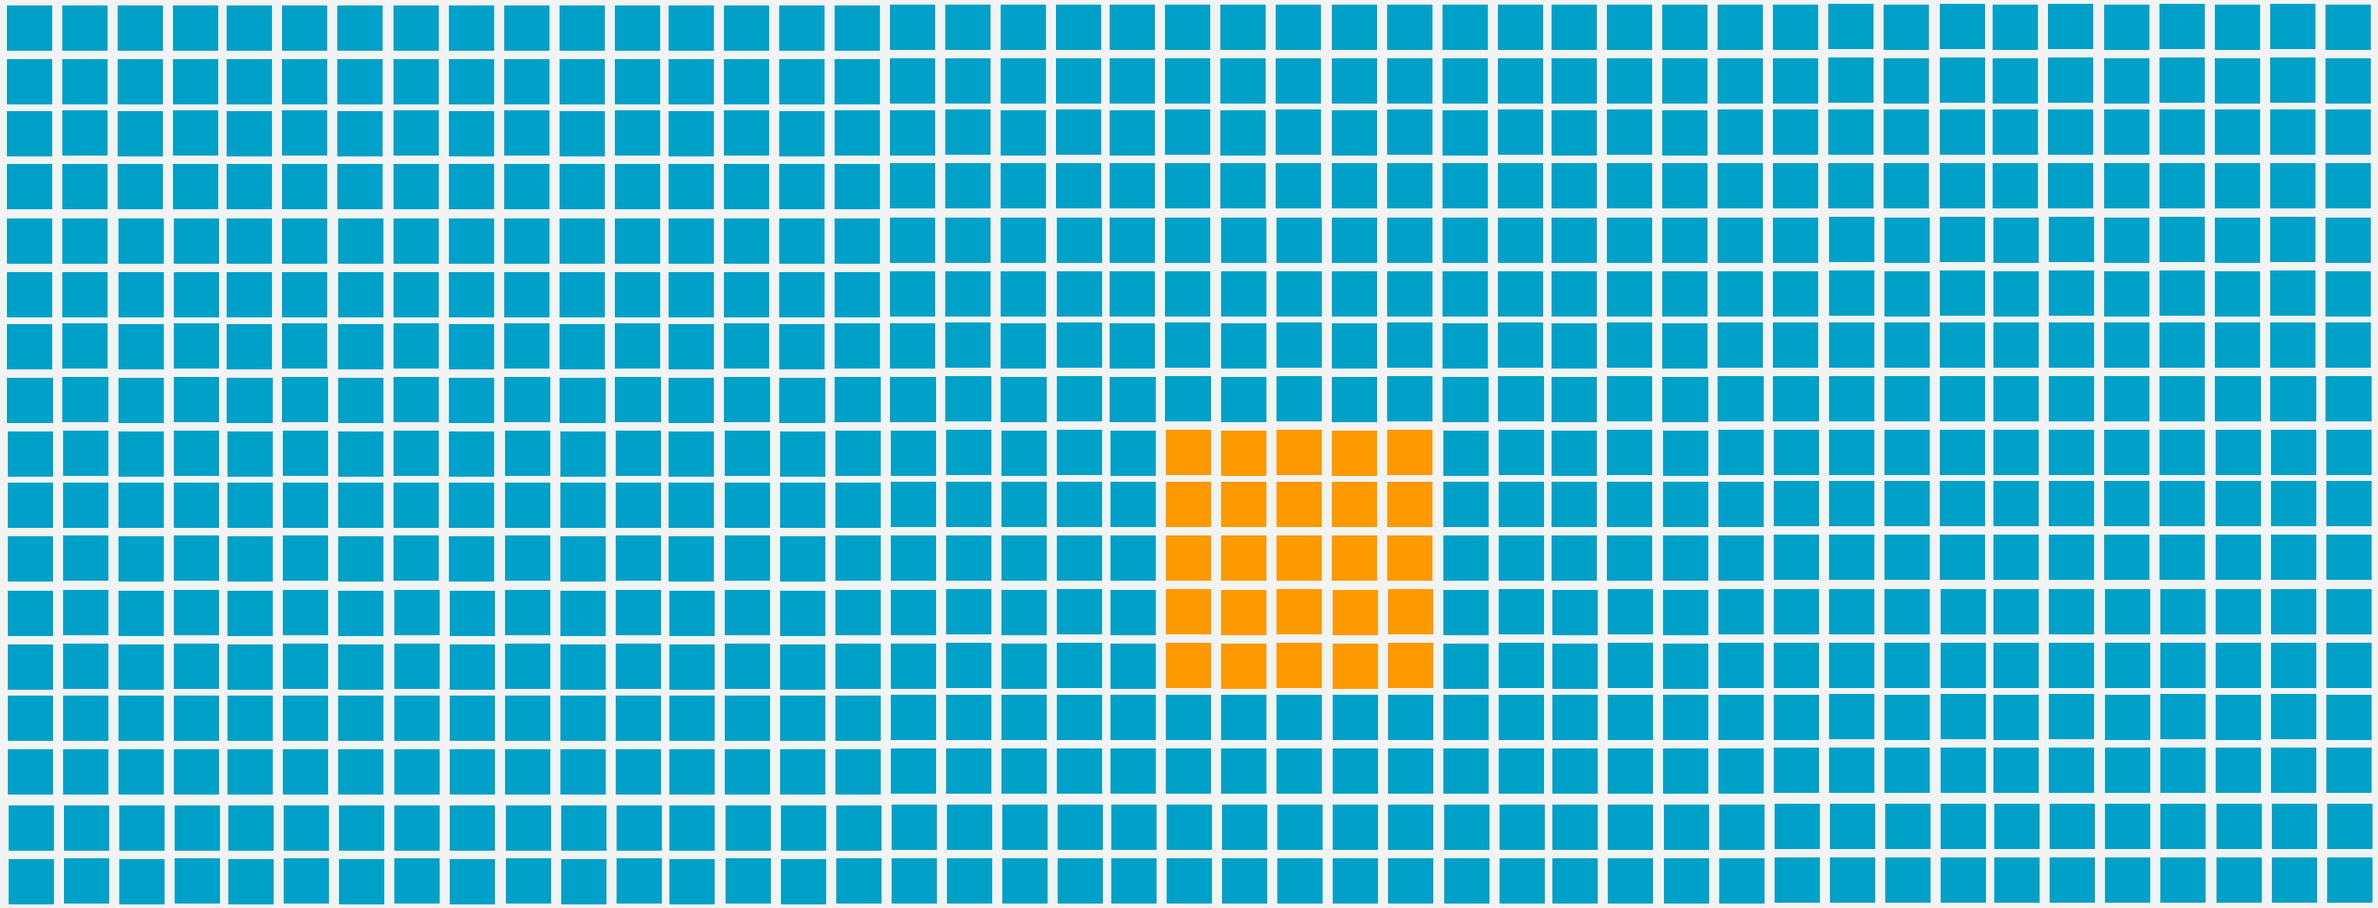

aws

Power Laws

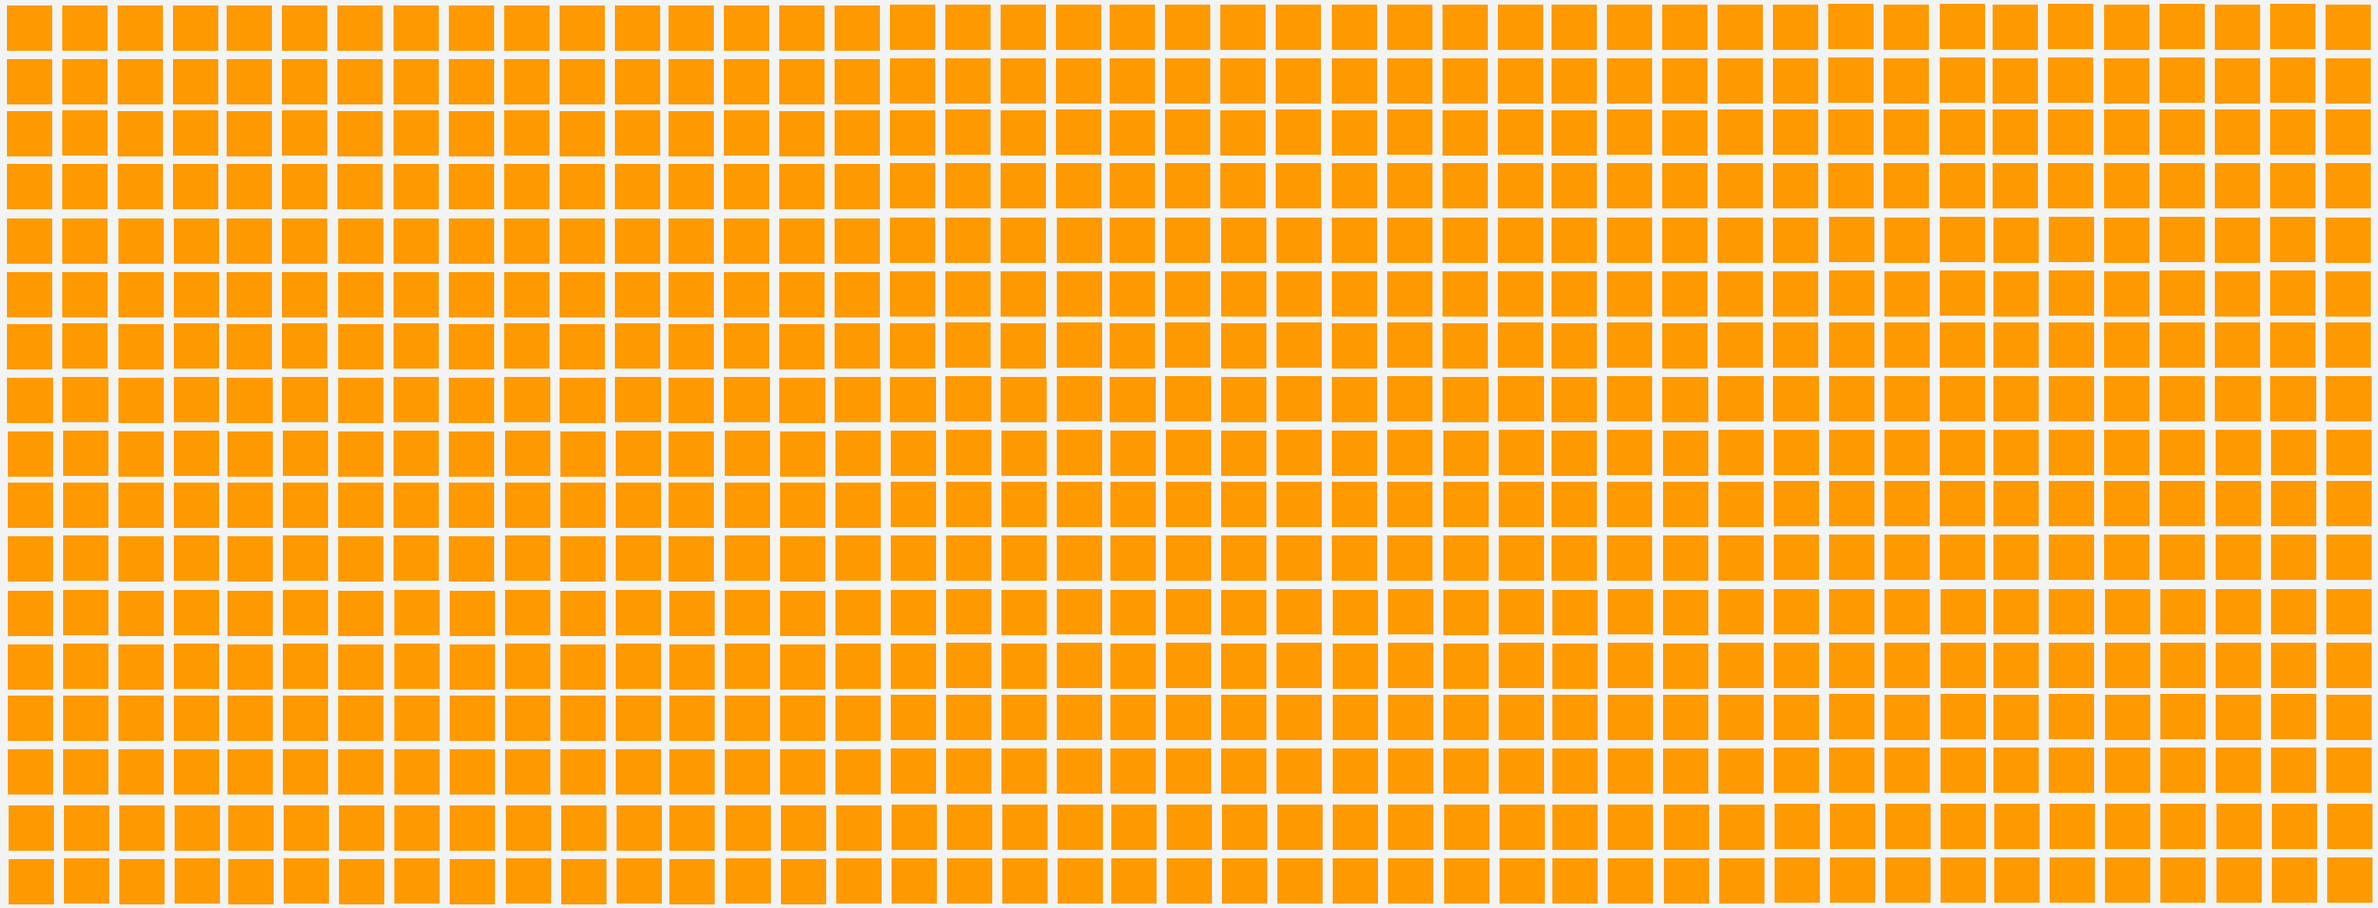

aws

Power Laws

- First, compartmentalize.
- More compartments means relatively smaller blast radius.
- Many real-world control systems reflect this lesson of scale.
- What next?

Power Laws

- Exponential Back-off
    - Brings our own power-law to the table
- Rate-limiters
    - Simple token buckets can be incredibly effective
- Working Backpressure
    - AWS SDK retry strategy = Token buckets + Rate-Limiters + persistent state

aws

Power Laws

AWS Architecture Blog

Exponential Backoff And Jitter

by Marc Brooker | on 04 MAR 2015 | in Architecture | Permalink | Share

Introducing OCC

Optimistic concurrency control (OCC) is a time-honored way for multiple writers to safely modify a single object without losing writes. OCC has three nice properties: it will always make progress as long as the underlying store is available, it's easy to understand, and it's easy to implement. DynamoDB's conditional writes make OCC a natural fit for DynamoDB users, and it's natively supported by the DynamoDBMapper client.

While OCC is guaranteed to make progress, it can still perform quite poorly under high contention. The simplest of these contention cases is when a whole lot of clients start at the same time, and try to update the same database row. With one client guaranteed to succeed every round, the time to complete all the updates grows linearly with contention.

aws

3

X-Ray Vision: Liveness and Lag

aws

X-Ray Vision: Liveness and Lag

- Operating on old information can be worse than operating on no information
- Simple example: system gets very busy and workflows and metrics pipelines can build up
- Ephemeral “shocks” such as spiky loads or brief outages can end up taking very long to recover

aws

X-Ray Vision: Liveness and Lag

- Strive for O(1) scaling as much as possible
- Provision everything, every time
- Report everything, every time
- Do everything, every time

aws

X-Ray Vision: Liveness and Lag

- If you need to use a bus or queue, think carefully about limits on the size of that queue

- In general: short queues are safer

- LIFO queues can be a great strategy for information channels
    - Naturally prioritizes recent state
    - Out of order back-fill for any “catching up”

aws

4

X-Ray Vision: False Functions

aws

False functions

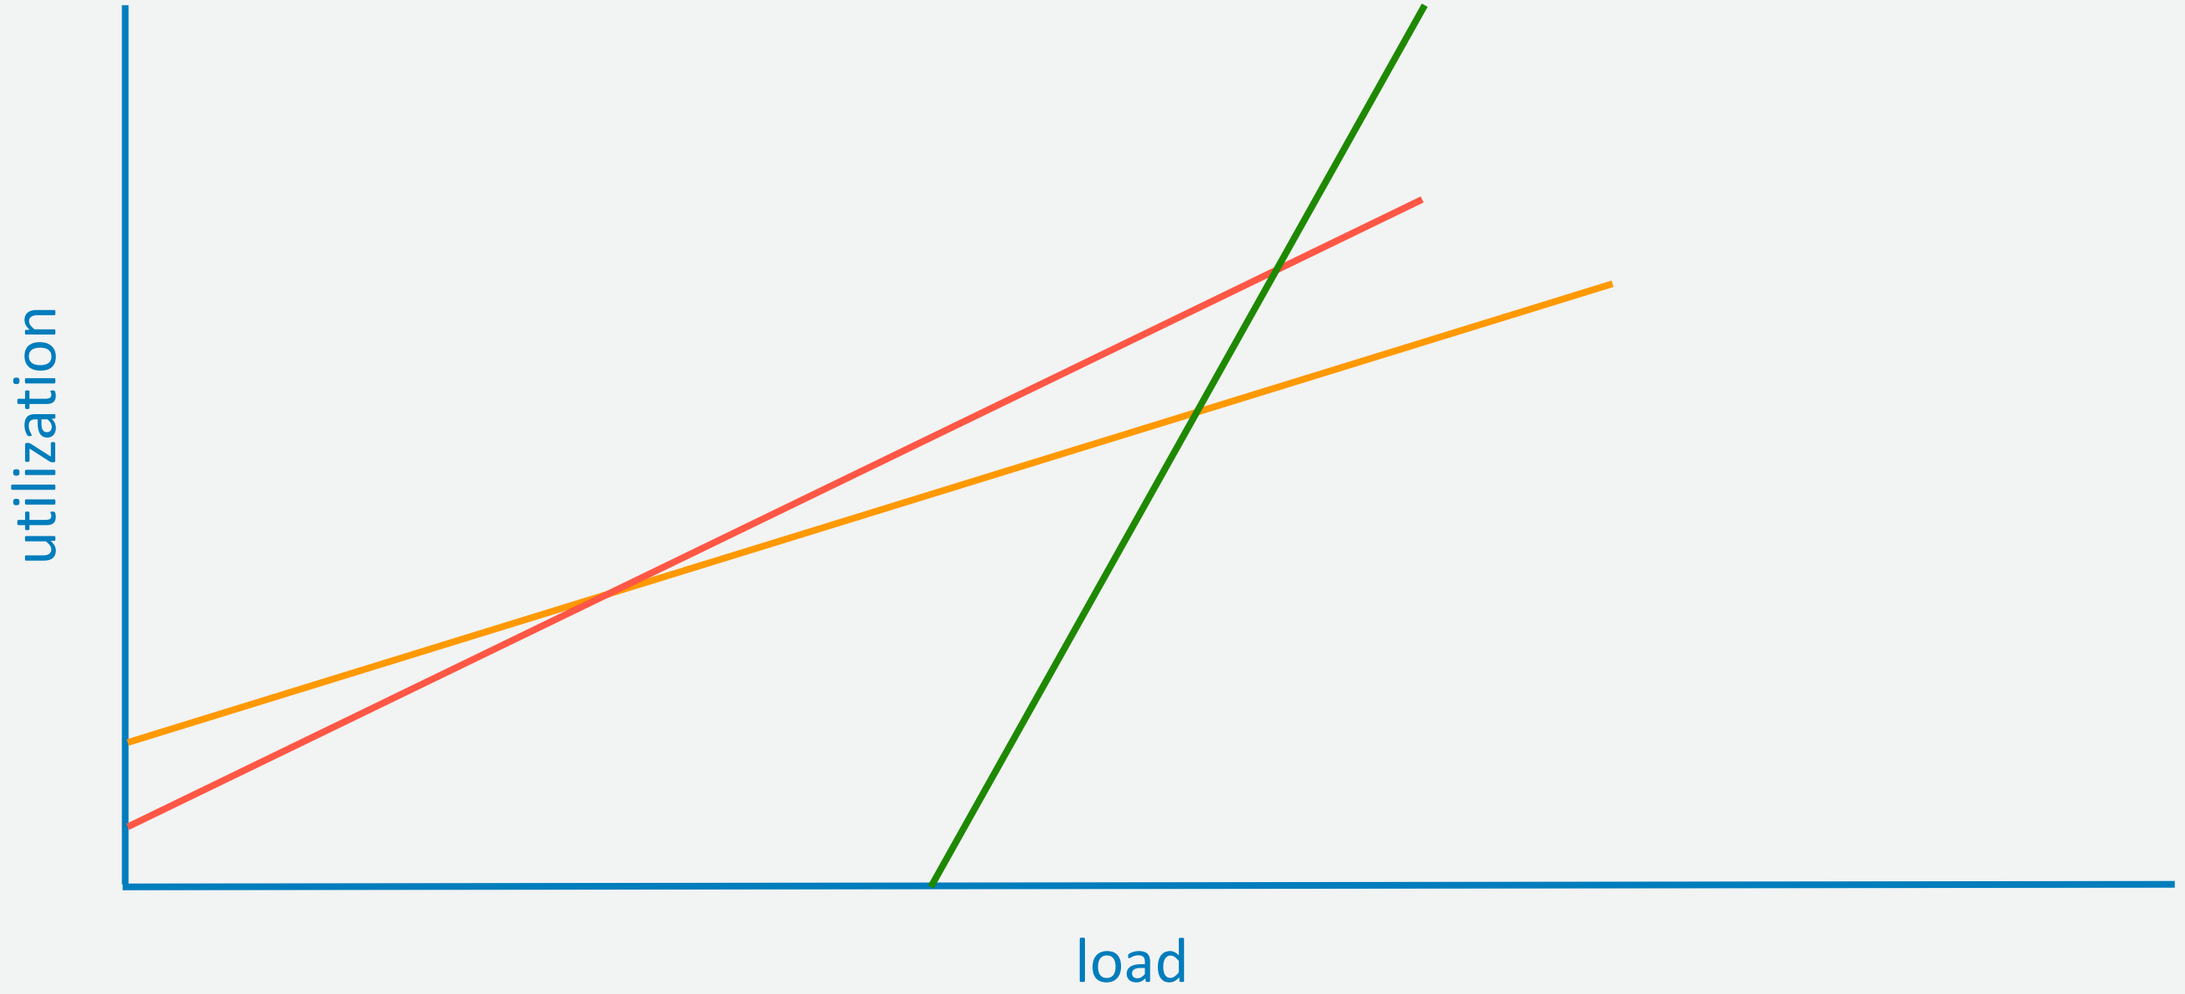

aws

False functions

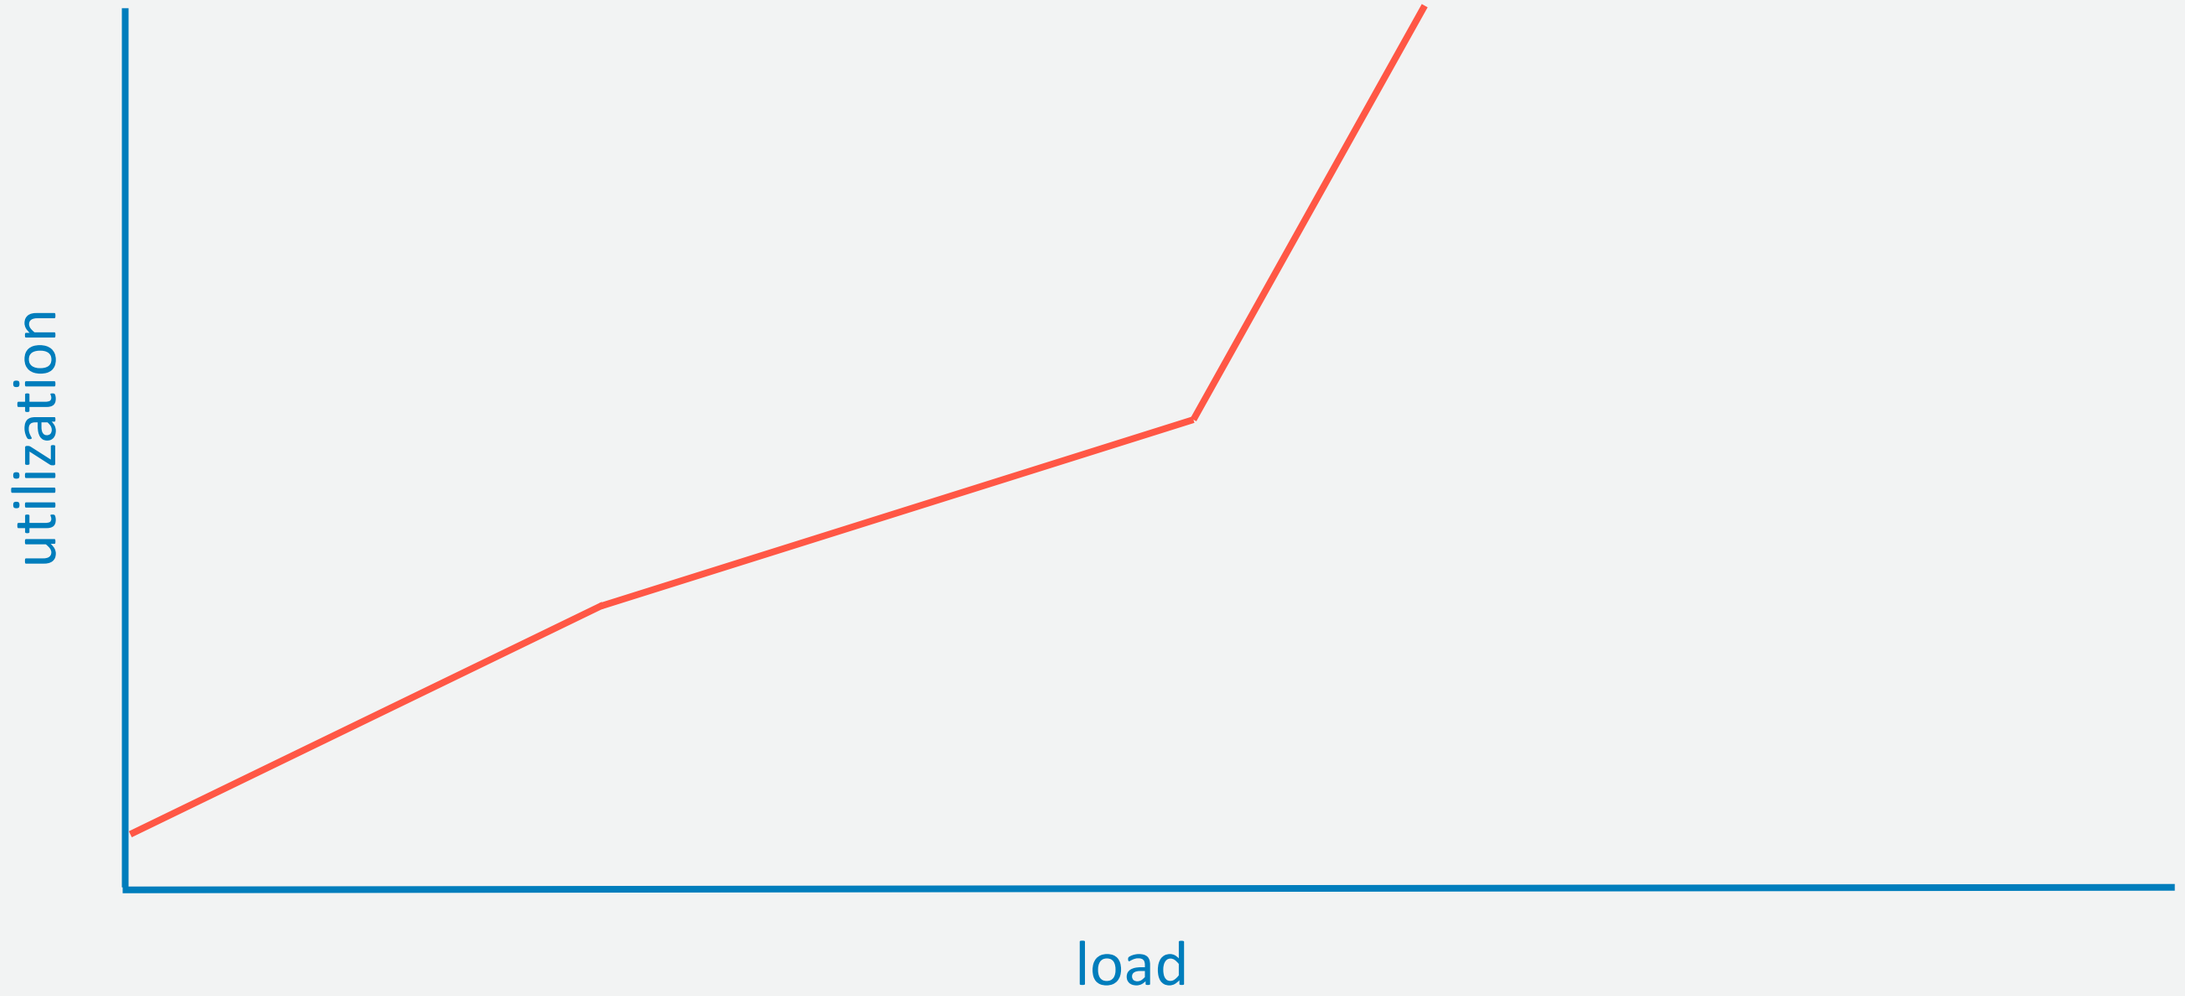

aws

False functions

- Hall of fame false function: Unix load
- Runners-up: system latency, network latency
- Hard to predict Garbage Collector behavior can be confounding
- CPU can be surprisingly effective

aws

5

X-Ray Vision: Edge Triggering

aws

Edge Triggering

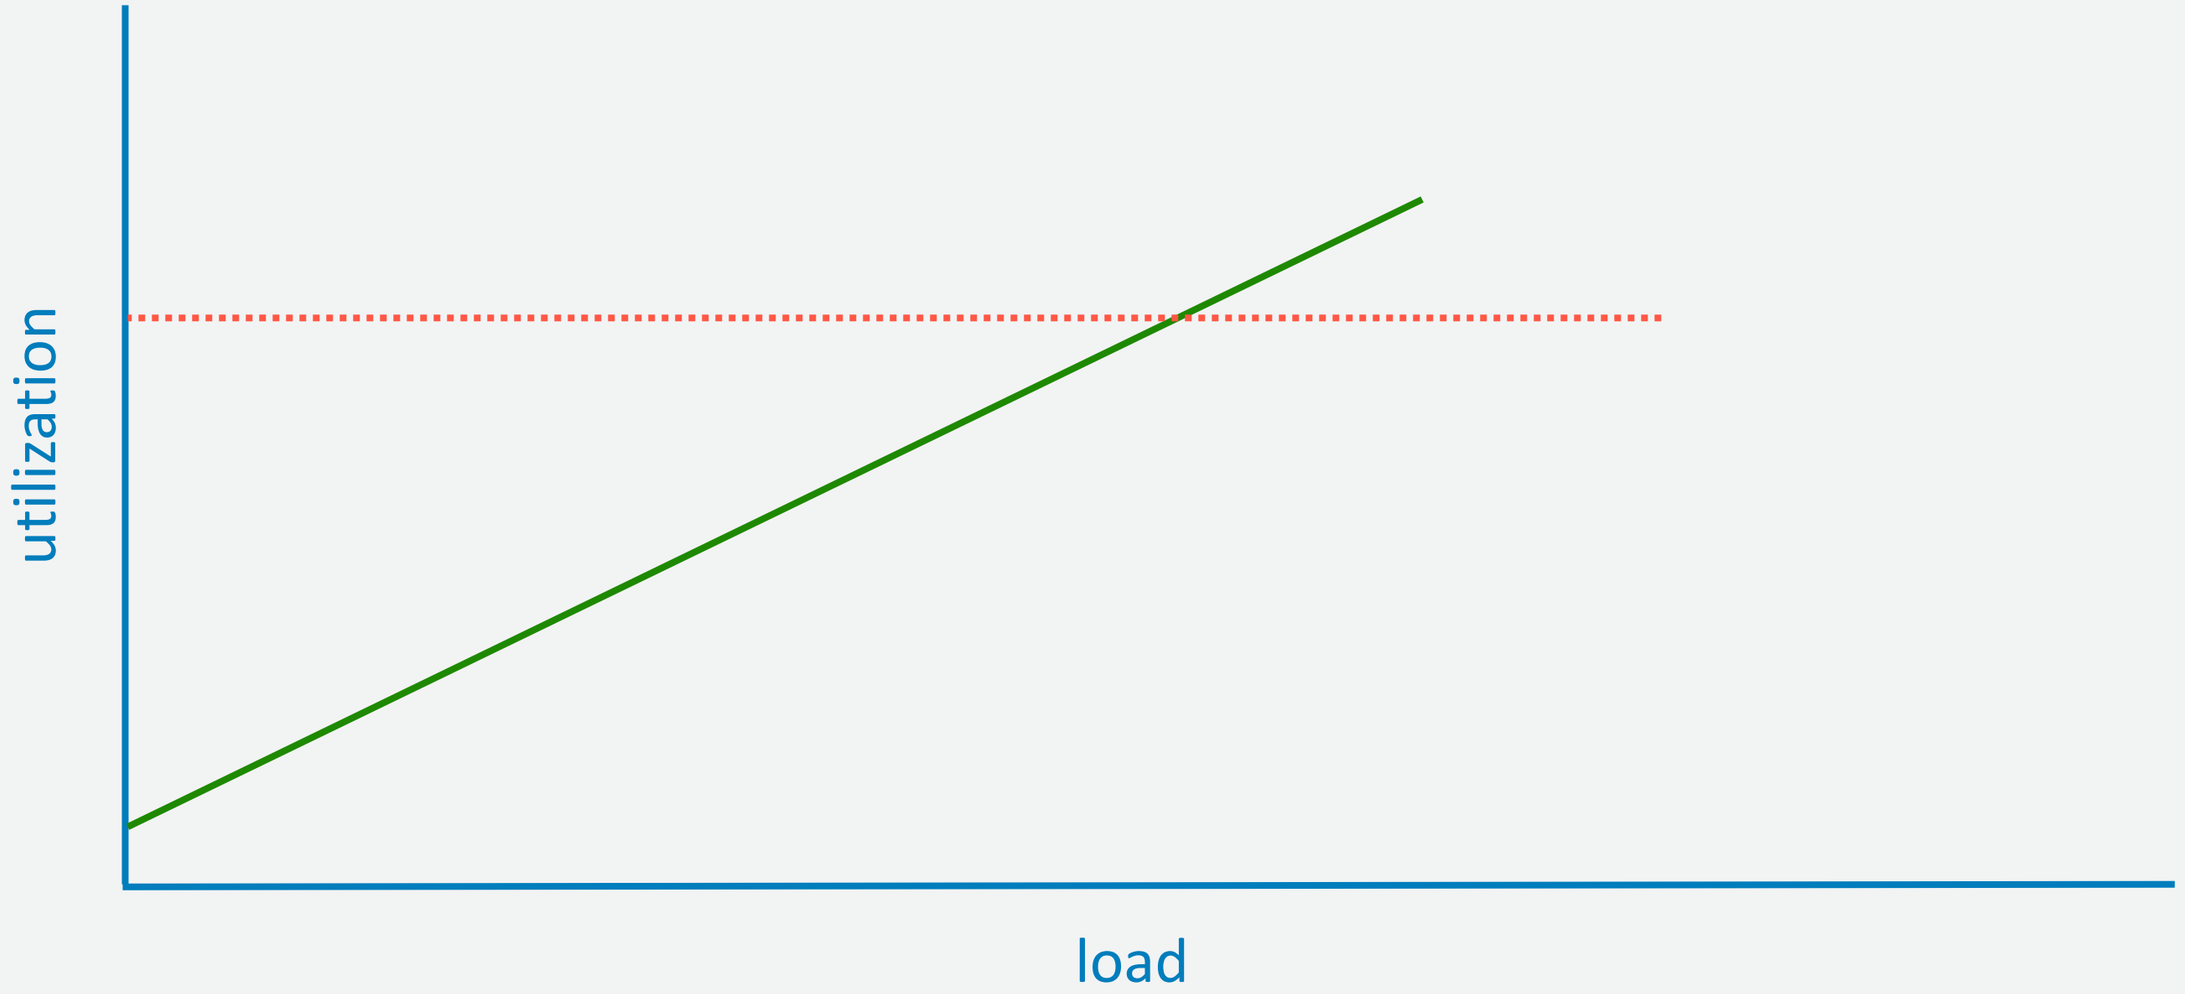

aws

Edge Triggering

- Edge Triggering invites modal behavior
- Often the new mode kicks in at a time of high-stress
- Edge Triggering often associated with the “Deliver exactly once” problem
- O.k. for alerting humans but usually an anti-pattern for control systems

aws

Summary

aws

Summary

- “Measure first” and “Integrate feedback” are deeply rewarding concepts
- Right now, this knowledge is highly leveraged
- We can think of distributed systems in terms of control theory, with 100 years of powerful mental models available
- Control Theory can help us formally analyze the stability of systems

aws

Q&A

Colm MacCárthaigh

aws

Thank you!