The SRE I aspire to be

Yaniv Aknin // @aknin


Google Cloud

#SREcon Dublin 2019

Who is this guy

- **Google SRE since 2013**
  Most recently *GCP's Quantitative Reliability Lead*
- **Jack of all trades**
  Equal parts SRE, dev, and /pro(duct|ject) manager/
- **Opinions my own**
  But I owe a lot here to others

Who is this guy

- **Google SRE since 2013**
  Most recently *GCP's Quantitative Reliability Lead*
- **Jack of all trades**
  Equal parts SRE*, dev, and /pro(duct|ject) manager/
- **Opinions my own**
  But I owe a lot here to others

*NB: what does "SRE" really mean?

Wikipedia says Engineering is "**using scientific principles to design and build $THINGS**"

https://en.wikipedia.org/wiki/Engineering

Wikipedia says Engineering is "**using scientific principles to design and build `$THINGS`**"

https://en.wikipedia.org/wiki/Engineering

Imagine `THINGS="Reliability"`... how do we apply science to that?

> When you can measure what you are speaking about, and express it in numbers, you know something about it; but **when you cannot measure it, your knowledge is of a meagre and unsatisfactory kind**.

William Thomson (Lord Kelvin)
President of the Royal Society
Lecture on "Electrical Units of Measurement"
Published in "Popular Lectures", Vol. 1, 1883
(abridged to fit slide)

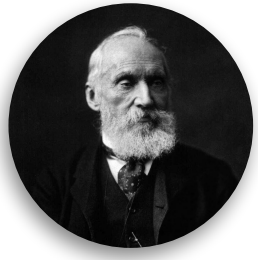

Measurably optimise reliability vs cost

On ops, user harm, and tradeoffs

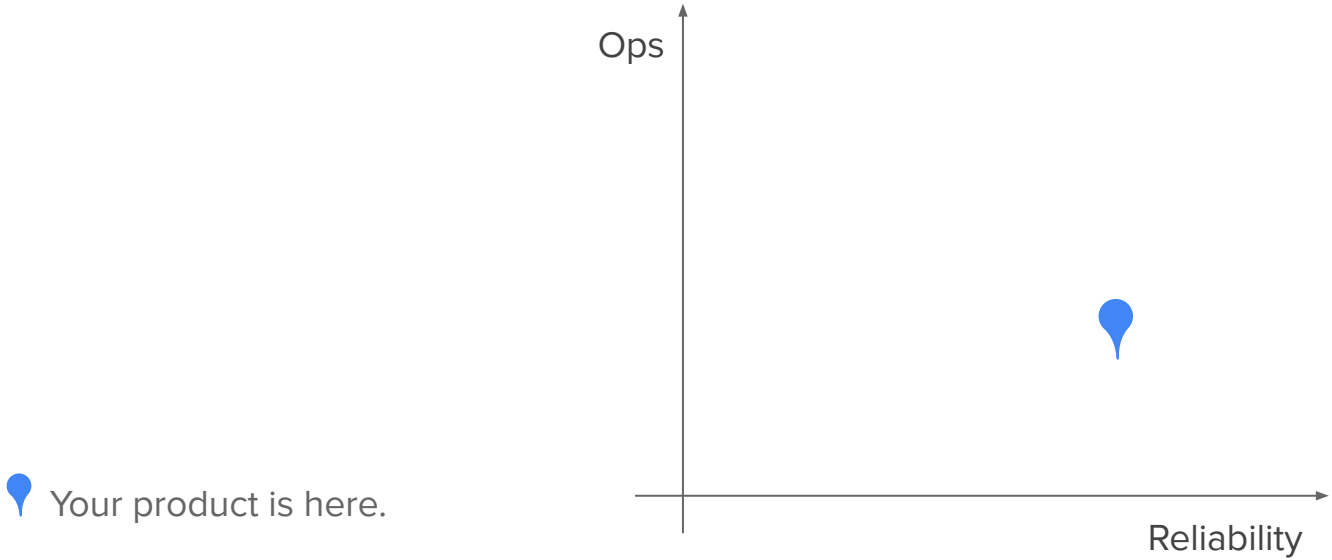

On ops, user harm, and tradeoffs

Your product is here.

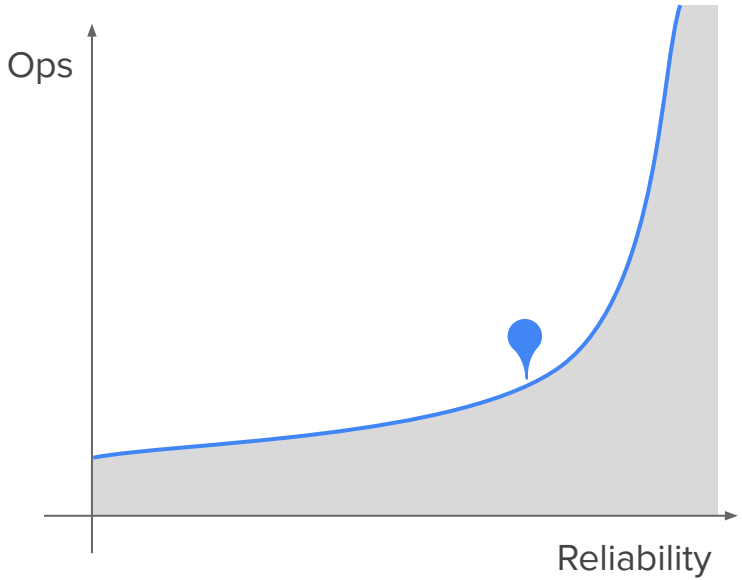

On ops, user harm, and tradeoffs

Your product is here.

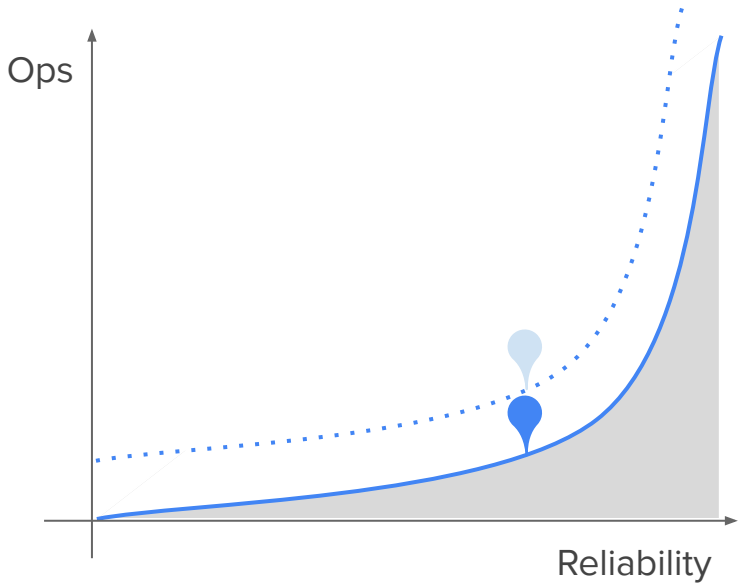

On ops, user harm, and tradeoffs

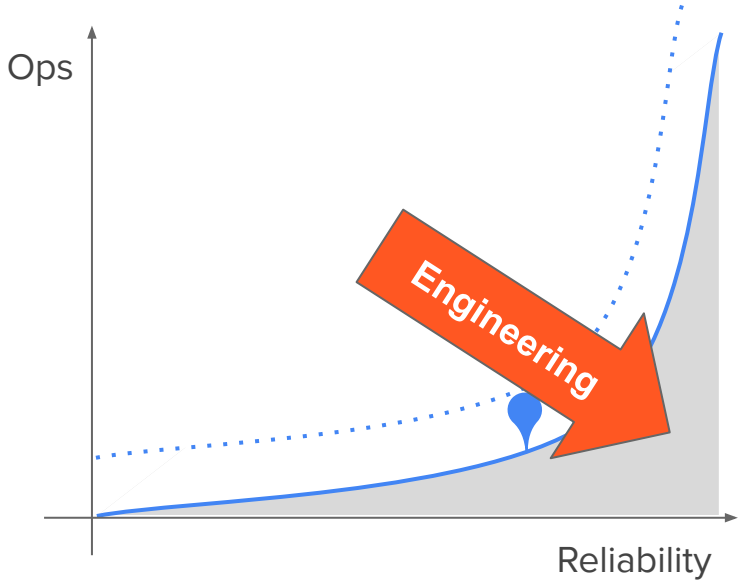

Your product is here.

Engineering reliability

A single HDD has an annualized failure rate (AFR) of ~1.5%

How can we build a more reliable logical disk?

Engineering reliability

A single HDD has an annualized failure rate (AFR) of ~1.5%

How can we build a more reliable logical disk?

RAID-1 (mirror) should theoretically offer $1.5\%^2$ (~0.02%) AFR

Assuming immediate disk replacement+replication after failure, completely independent disk failures, and no RAID related bugs. None of which is even remotely true, of course.

The (modest) reliability engineer toolbox

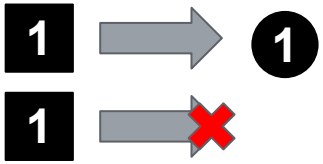

Redundant resource

Trade cost

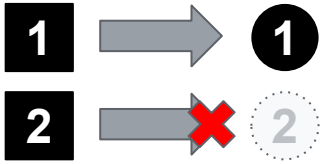

Degraded results

Trade quality

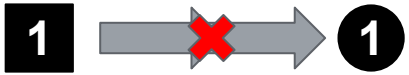

Retry transient failures

Trade latency

The (modest) reliability engineer toolbox

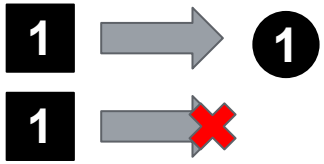

Redundant resource

Trade cost

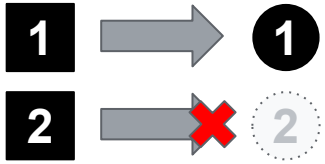

Degraded results

Trade quality

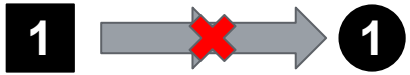

Retry transient failures

Trade latency

Any of these

Adds complexity

Compound/advanced reliability patterns

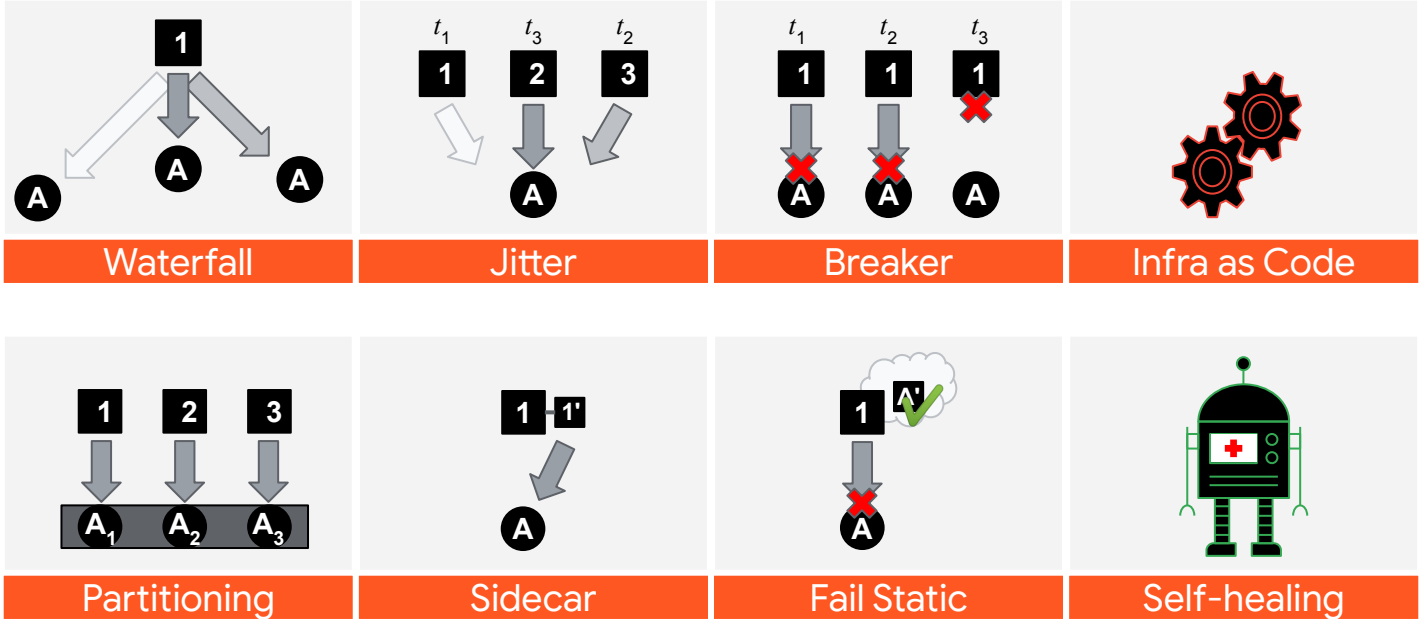


Innovation
(engineering, proactive, change)
Reliability
(support, reactive, preserve)

(support,
reactive,
preserve)

Reliability

(engineering,
proactive,
change)

Innovation

(engineering, proactive, change)

Reliability

(engineering, proactive, change)

Innovation

The Error Budget

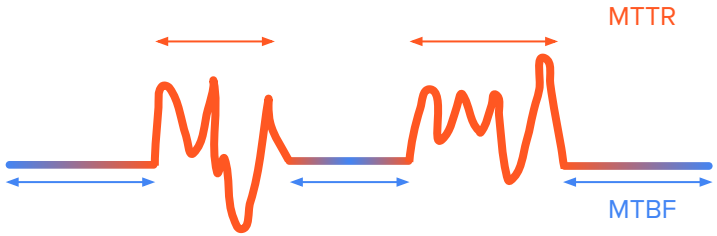

MTBF/MTTR

Challenge: fungible definition of "failure"

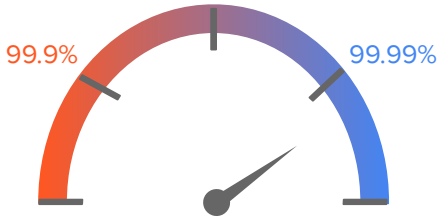

"9s" (e.g. "99.95% uptime")

Challenge: aggregating individual events into business credible 9s

You need "better quality" 9s!

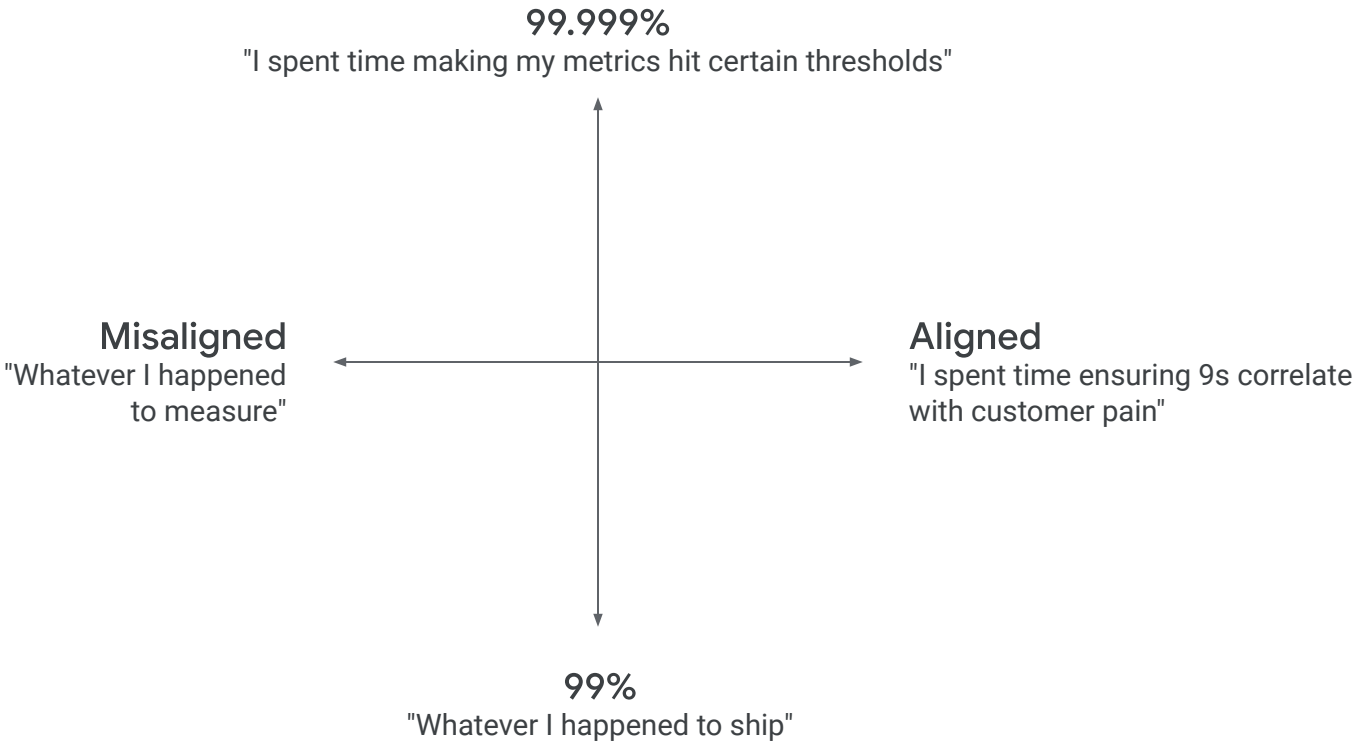

First move right, then move up

Measurably optimise reliability vs cost

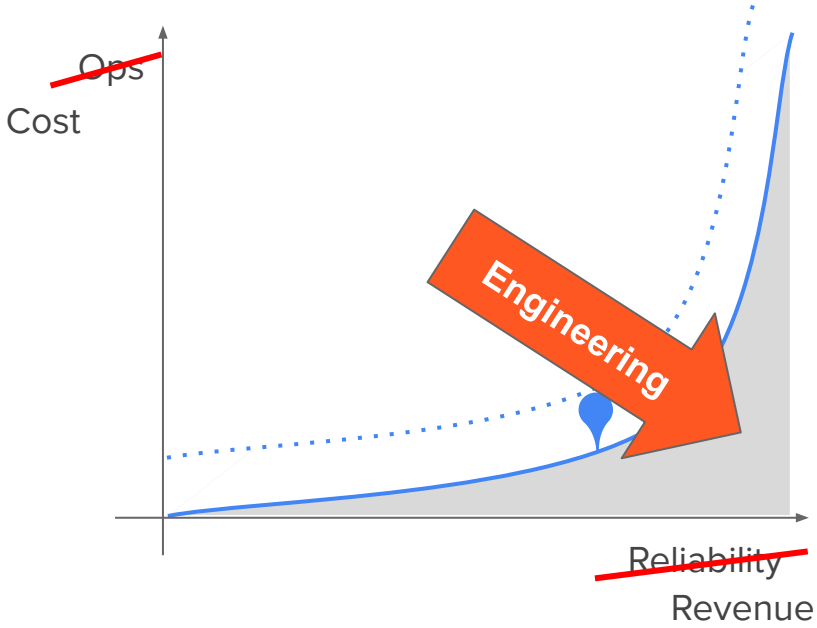

Your product is here.

Why is this hard?

- Scope
- Difficulty
- Cost++
- Misconceptions

Why is this hard? And why is it good?

- Scope
- Difficulty
- Cost++
- Misconceptions

- Leverage
- Cost--
- Precision

SRE team: a recipe

Fundamental

Monitoring
Alerting
Capacity planning
CI/CD & Rollouts
Load Balancing

SRE team: a recipe

Fundamental

Monitoring
Alerting
Capacity planning
CI/CD & Rollouts
Load Balancing

Advanced

System Architecture
Distributed Algorithms
Networking
Operating Systems

SRE team: a recipe

Fundamental

Monitoring
Alerting
Capacity planning
CI/CD & Rollouts
Load Balancing

Advanced

System Architecture
Distributed Algorithms
Networking
Operating Systems

Pioneering

Product Management
Data Science
Business Acumen
UX Research

The SRE I aspire to be

- Have a measurement of reliability
- The measurement is tied to project priorities
- Your ops work is tied to the measurement

The SRE I* aspire to be

- Have a measurement of reliability
- The measurement is tied to project priorities
- Your ops work is tied to the measurement

* Please remember this is *my* aspiration... tell me yours?

Thank you!

Yaniv Aknin // @aknin

Art credits

"Lord Kelvin", Messrs. Dickinson, London, goo.gl/RHF61Z, [cropped]
"Complex looking chart", MIT SERG, http://strategic.mit.edu, [recoloured]
"Gears" and "Robot", Google AutoDraw, CC4, [recoloured and adapted]
Yin Yang, https://en.wikipedia.org/wiki/File:Yin_yang.svg [recoloured]