Heresy & Evangelism

Schism in the church of monitoring

Hi! 👋

- Community @ elastic
- Reach me here:
  - aaron.aldrich@elastic.co
  - @CrayZeigh
- Slides are here:
  - noti.st/crayzeigh
- This picture is amazing, come @ me.

A word from our sponsor

- We make:
  - Elasticsearch
  - Logstash
  - Kibana
  - Beats
  - Elastic APM (open tracing, ooo)
- We host:
  - Elastic Search Service
  - Site Search
  - App Search
- You can run it all where ever
- Open Source
- **We're hiring** (Fully Distributed, oooh, aaah)
- Talk to me later

Let's find out where we're at.

How many of you deal with monitoring as a job function?

How many of you touch monitoring in some way?

Uptime

Performance/Resource Utilization

Response time?

Why?

Things Fall Apart*

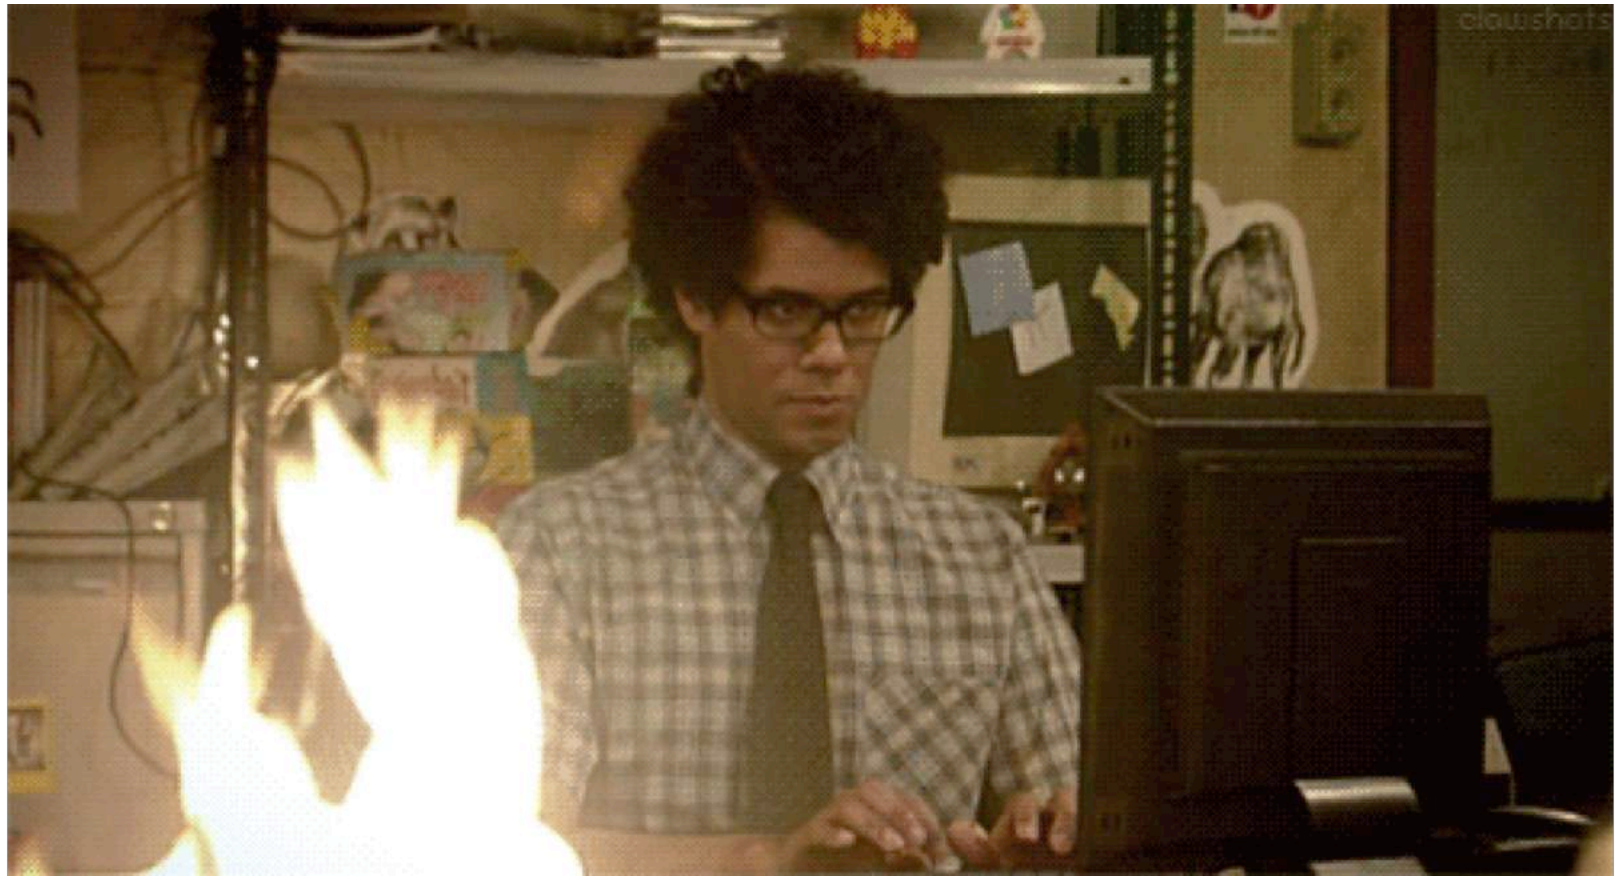

[*] something about a slouching beast

Incidents Suck

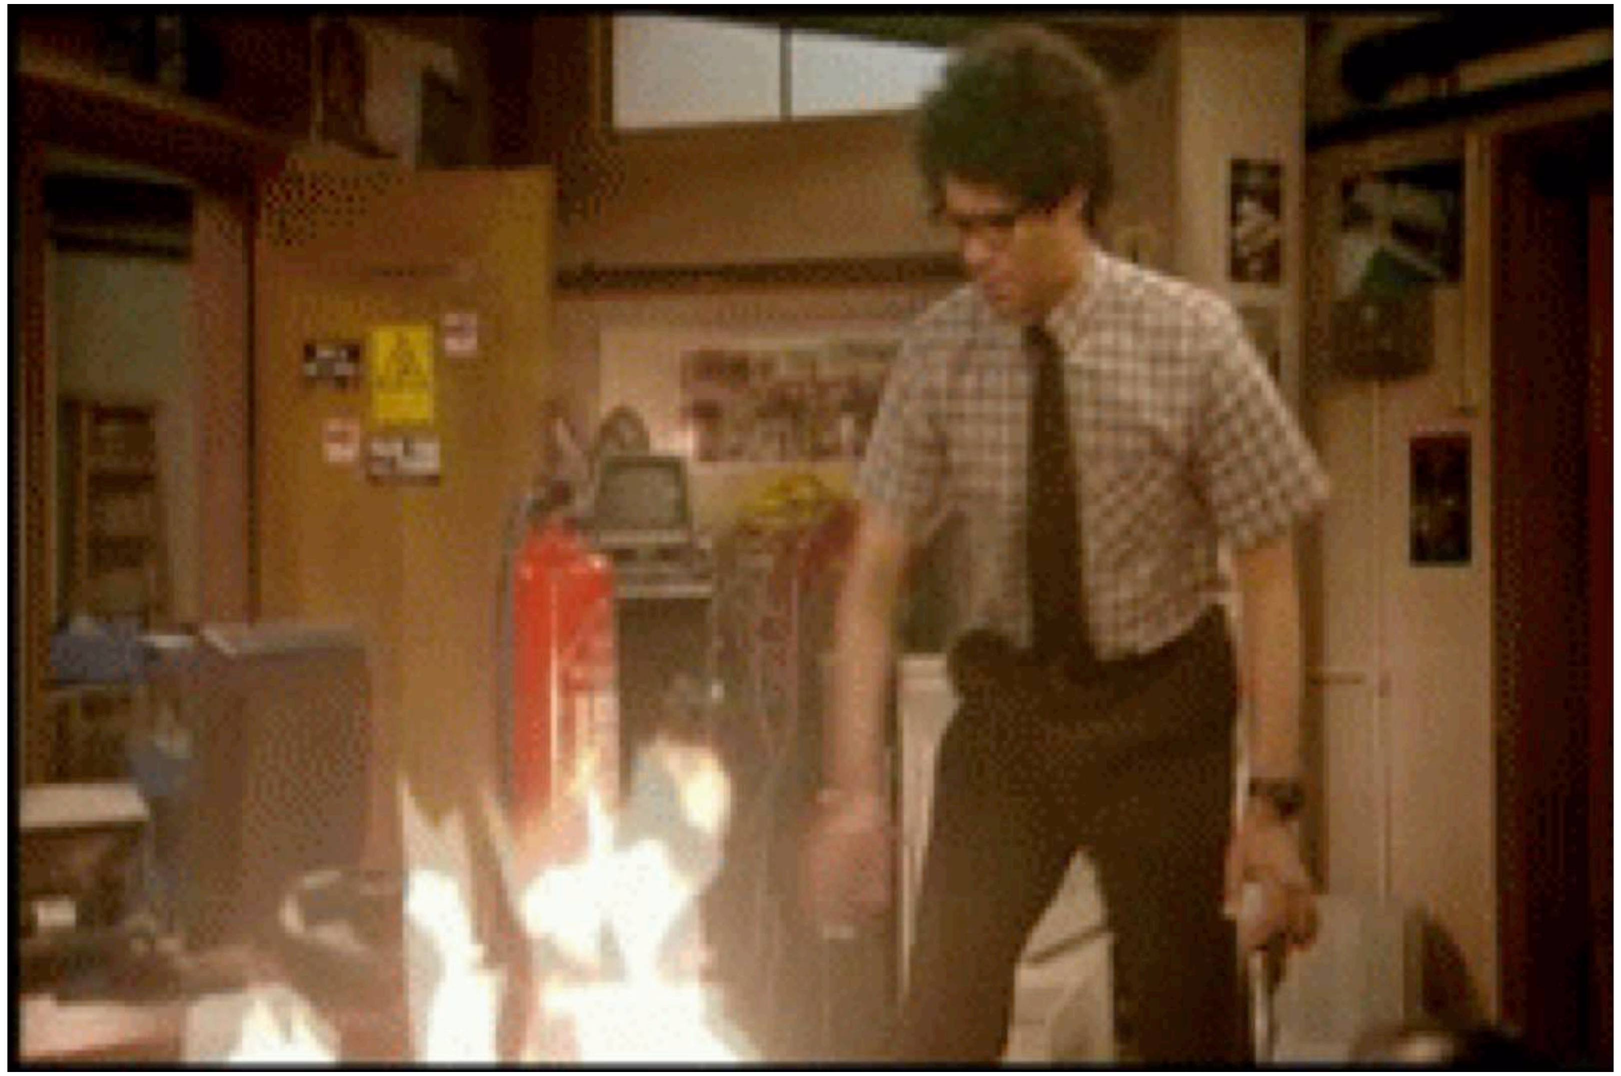

Locus of Control

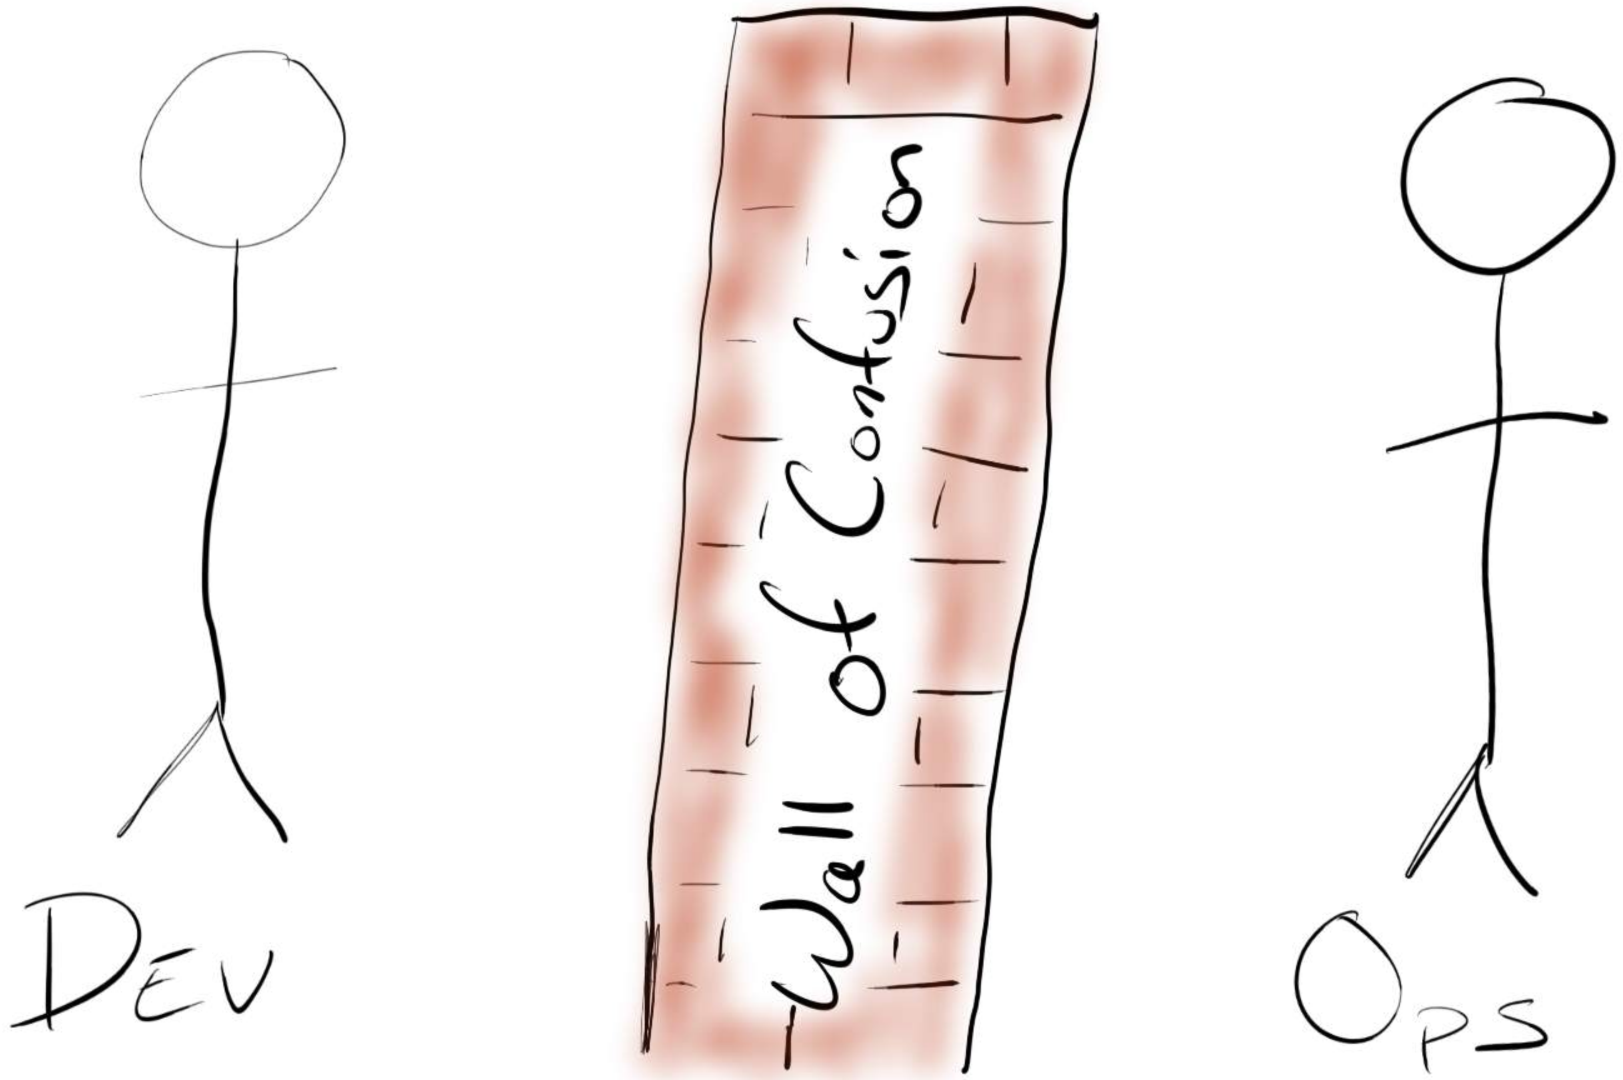
Dev
Wall of Confusion
Ops

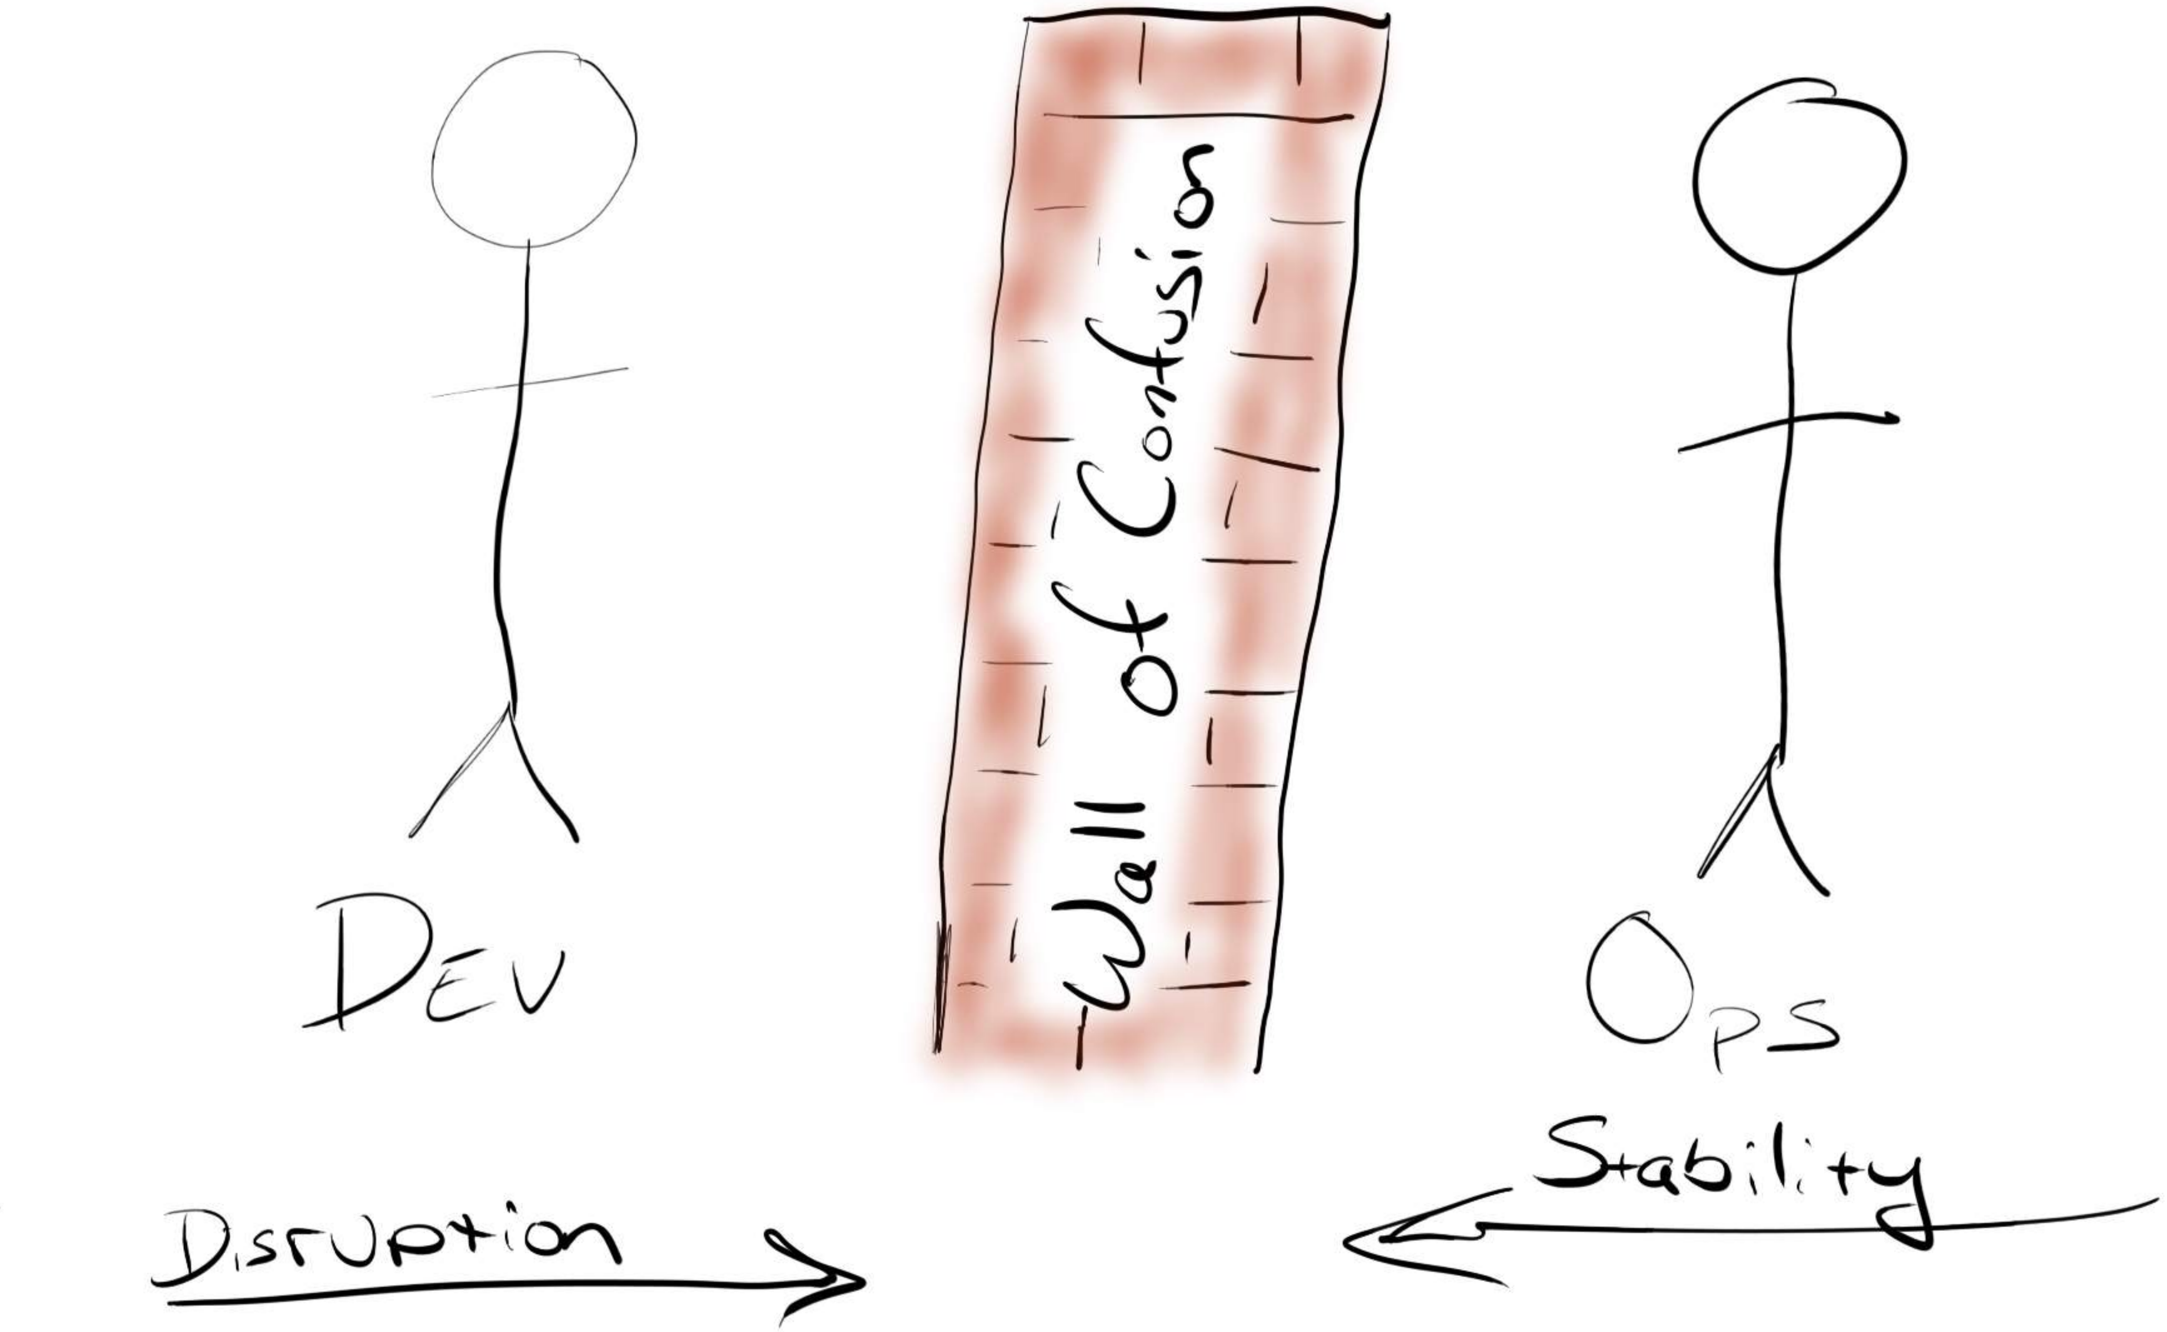
Dev
Wall of Confusion
Ops
Disruption
Stability

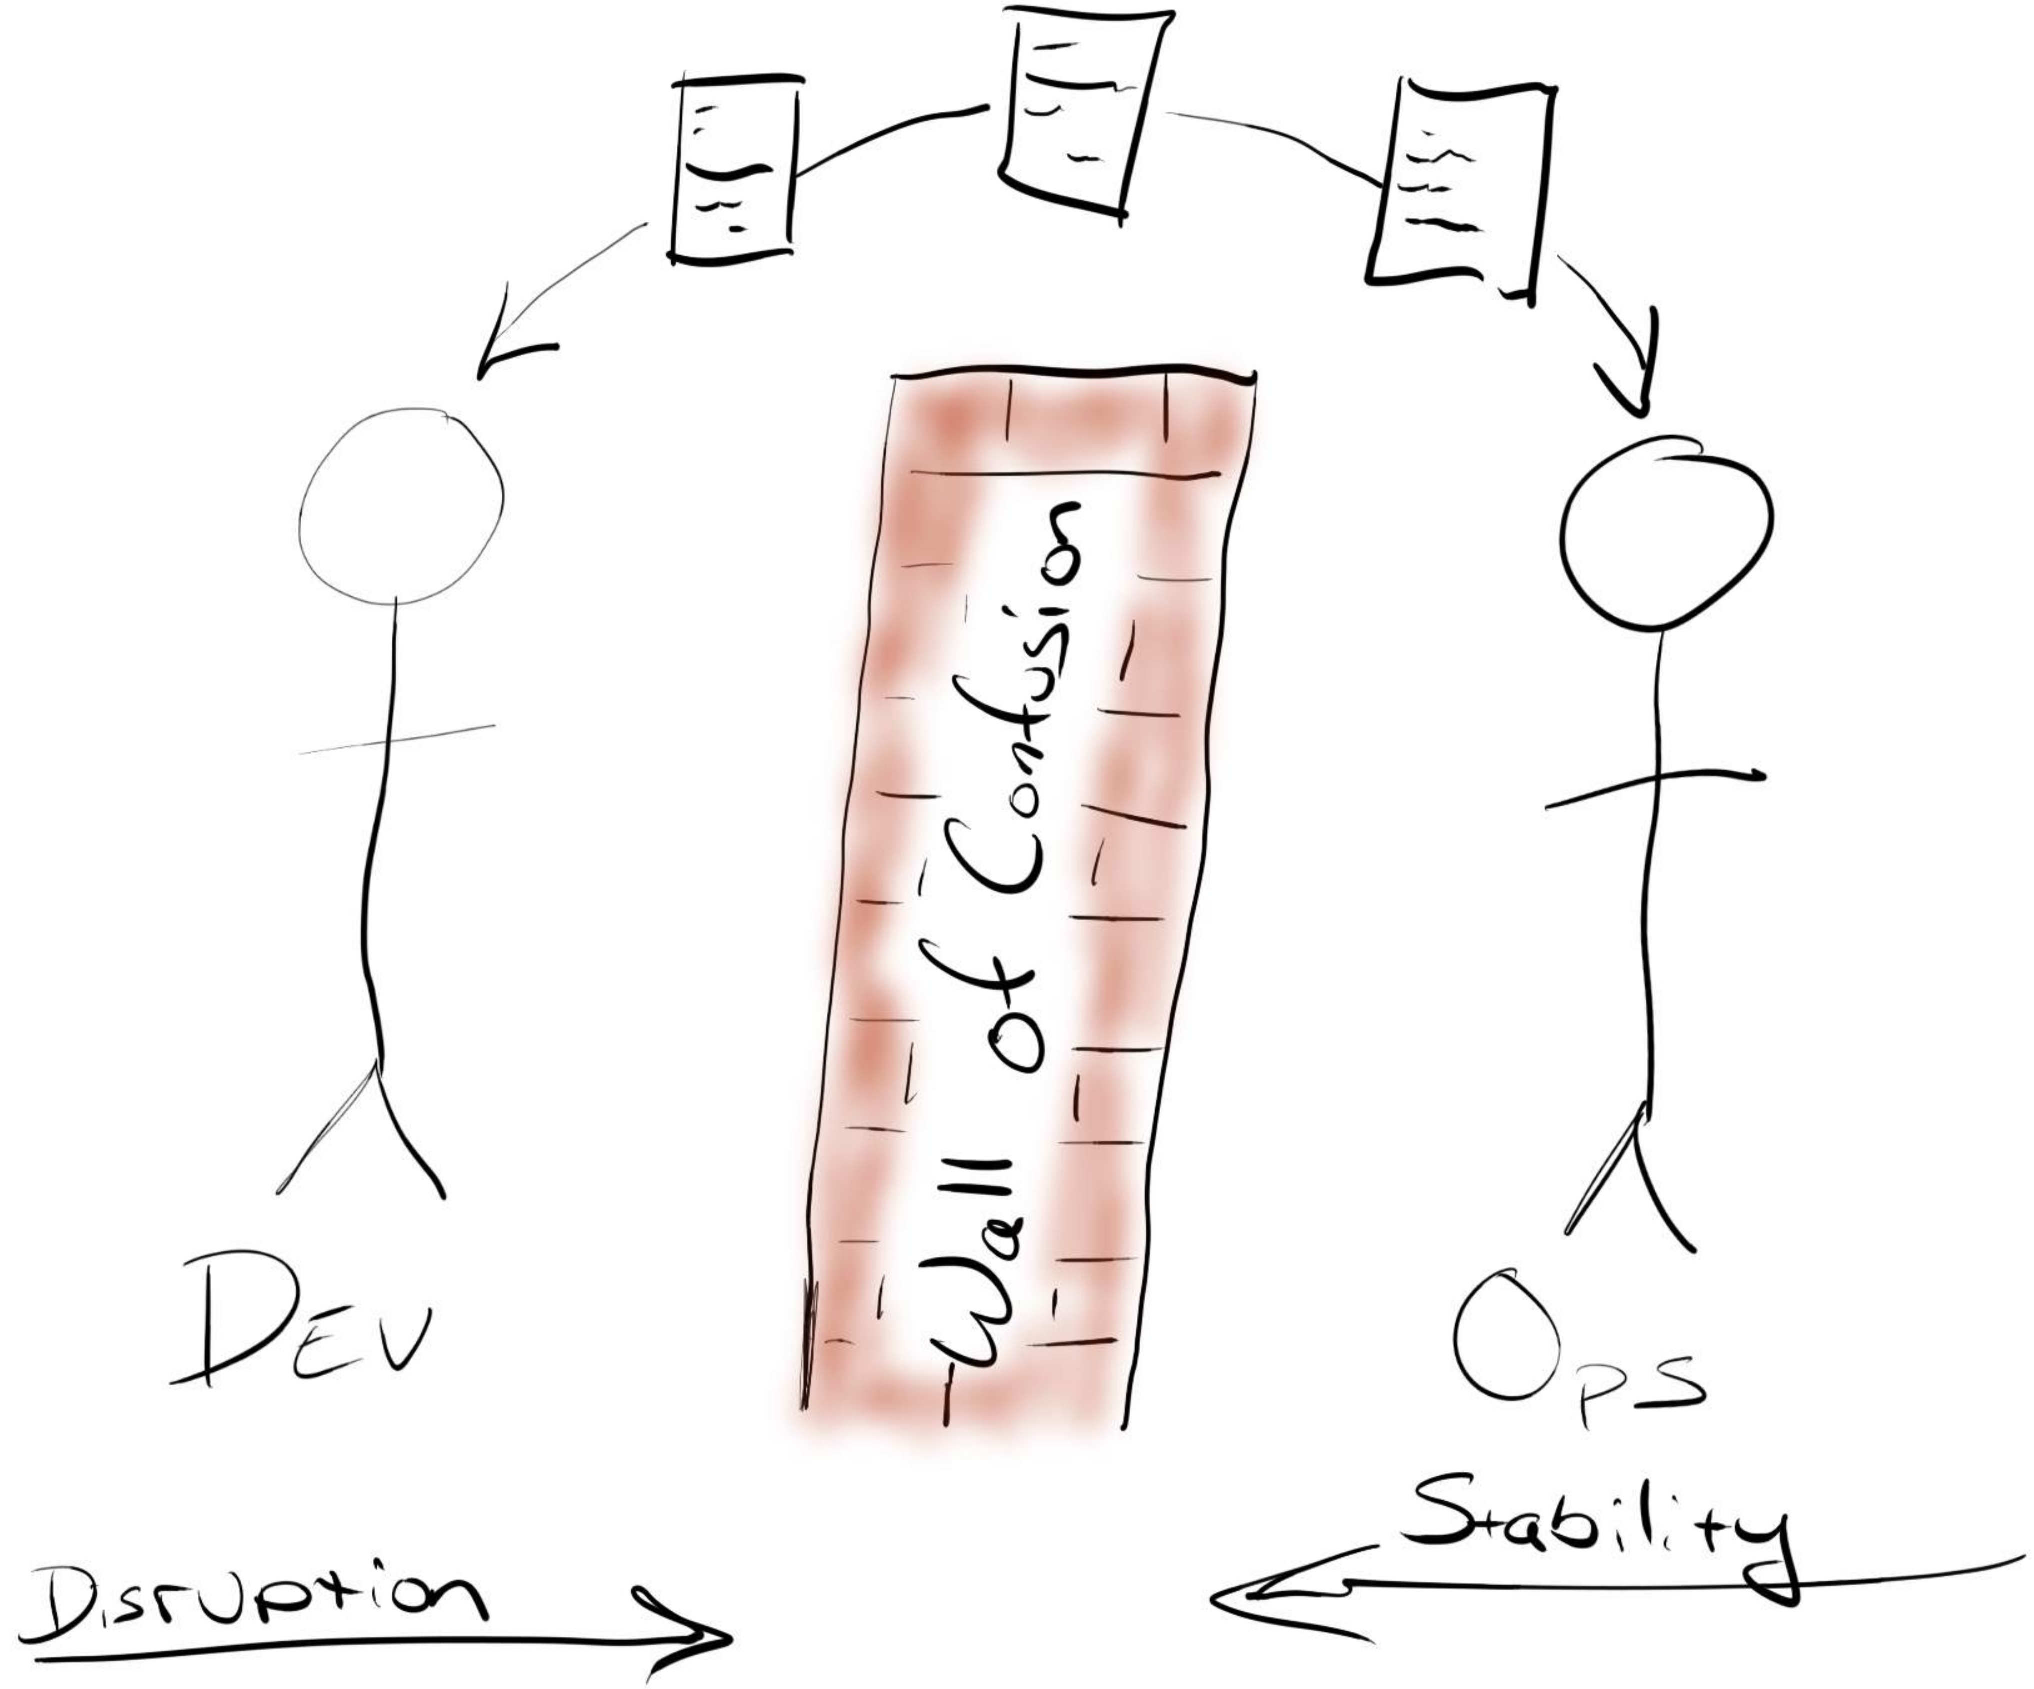
Dev
Wall of Confusion
Ops
Disruption
Stability

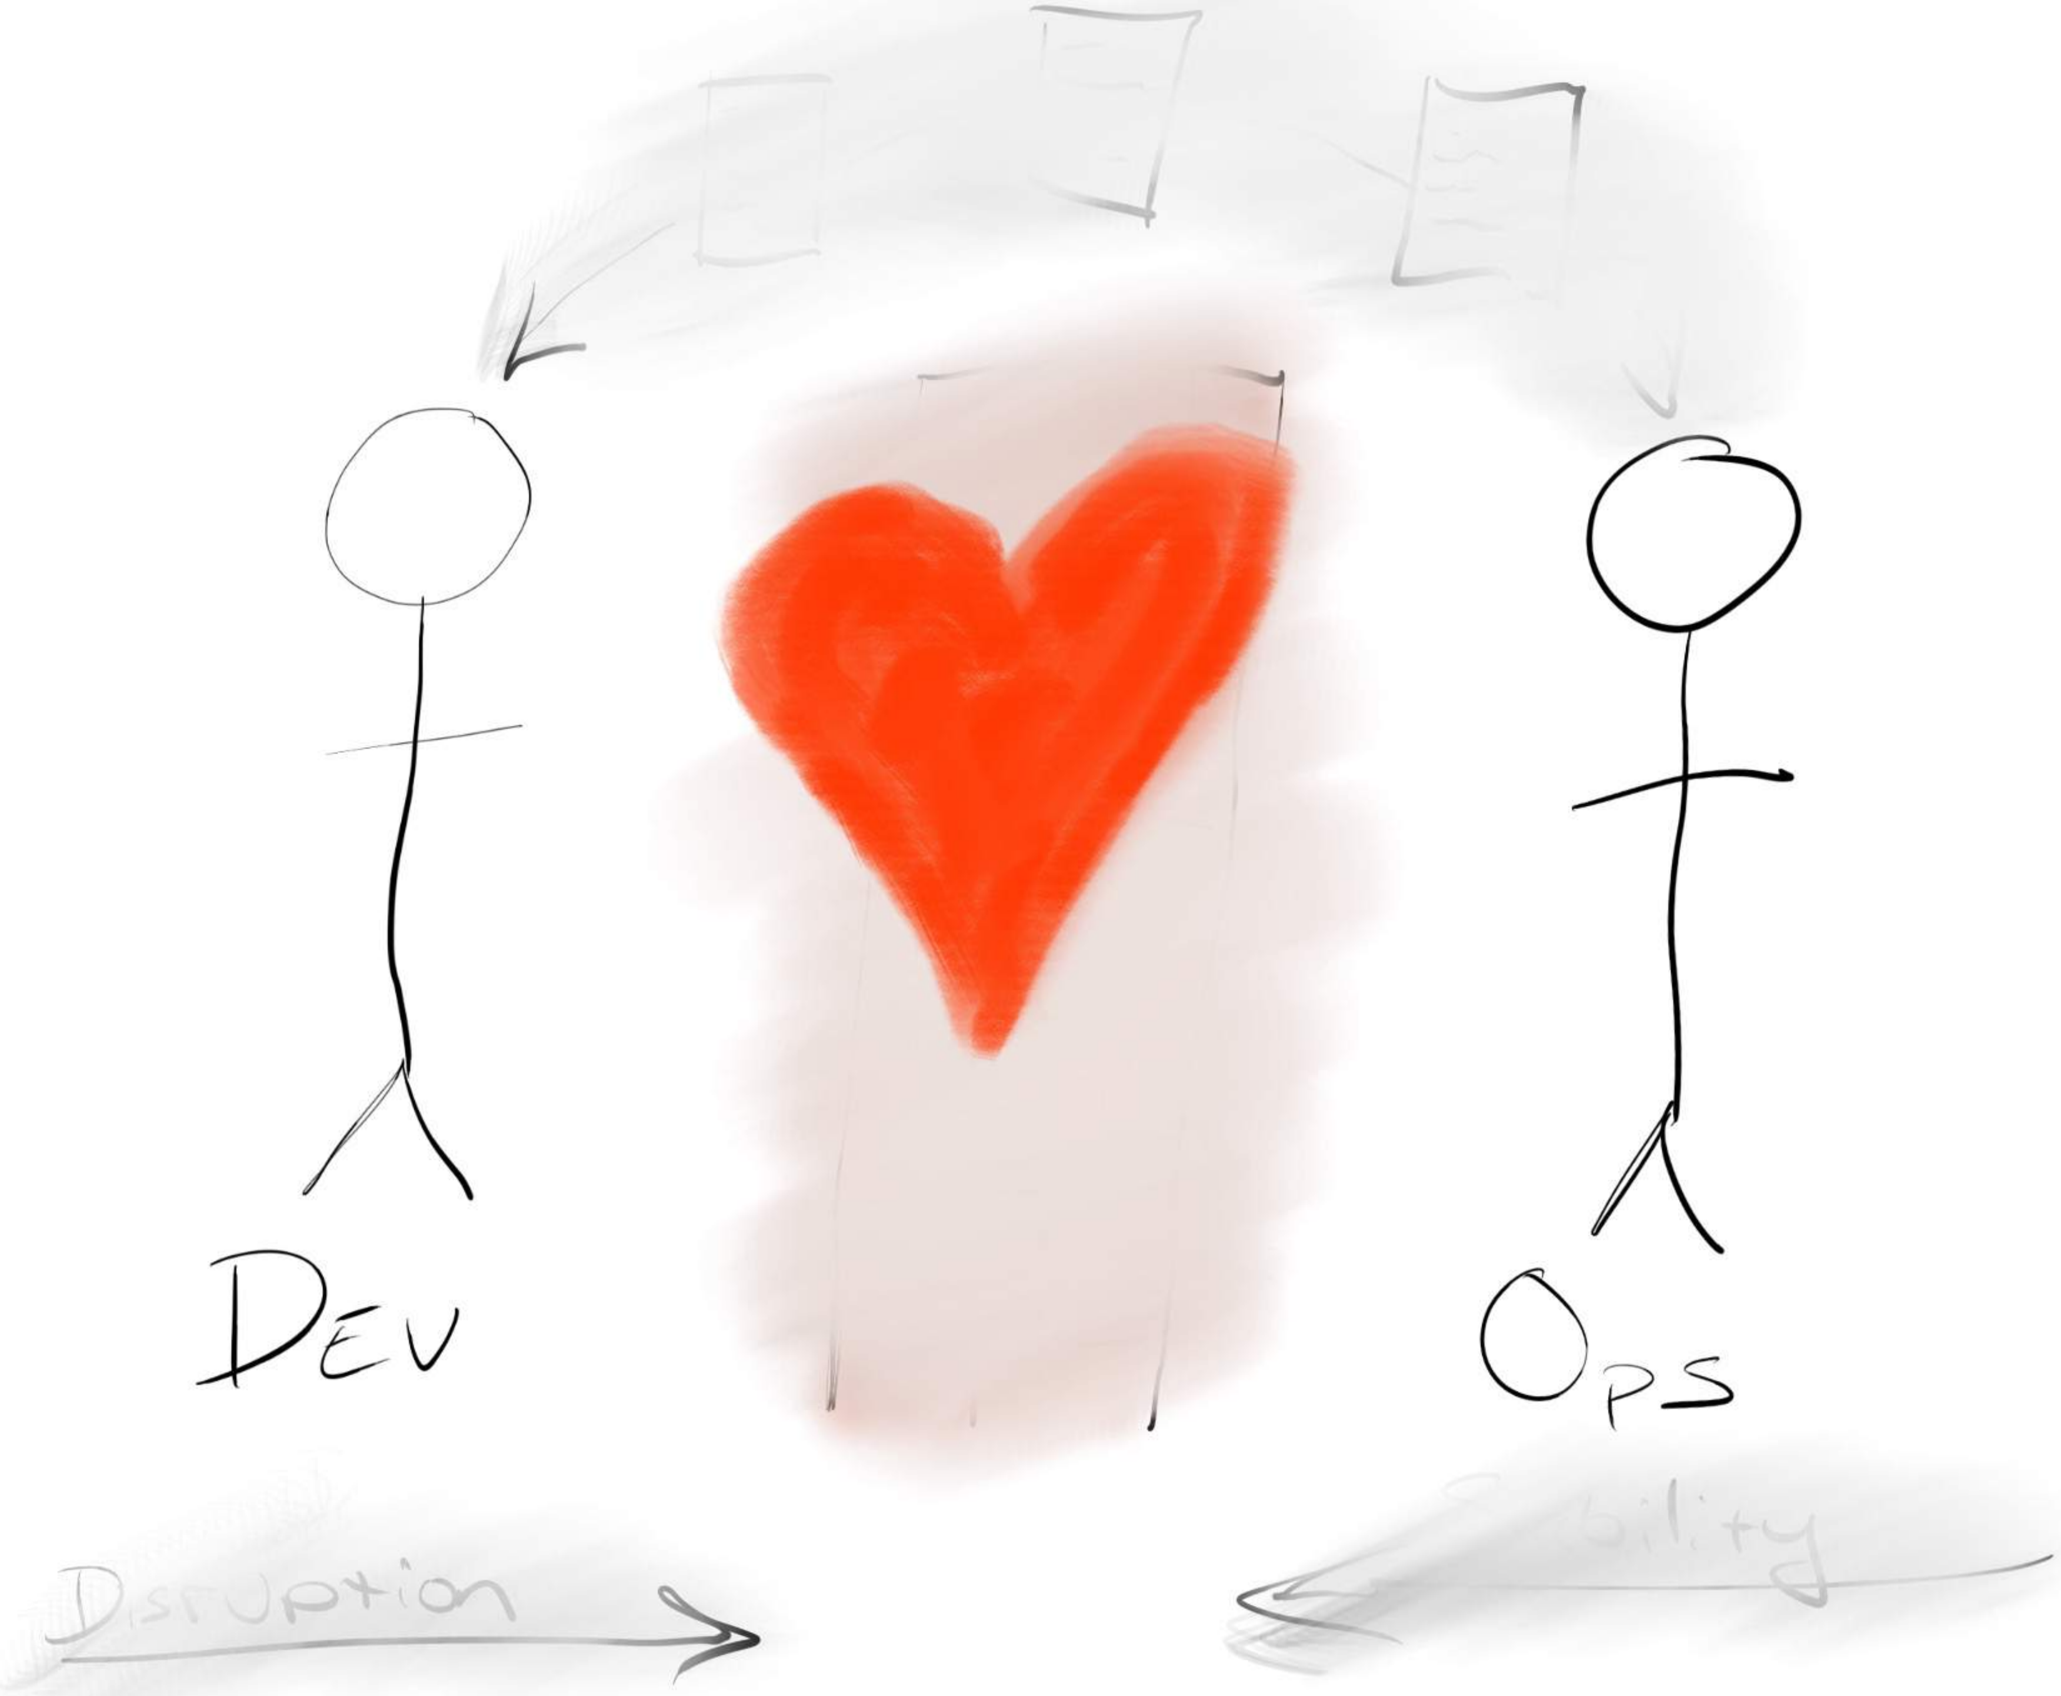
Dev
Ops
Disruption
Stability

100%

99.999%

Just a minute!

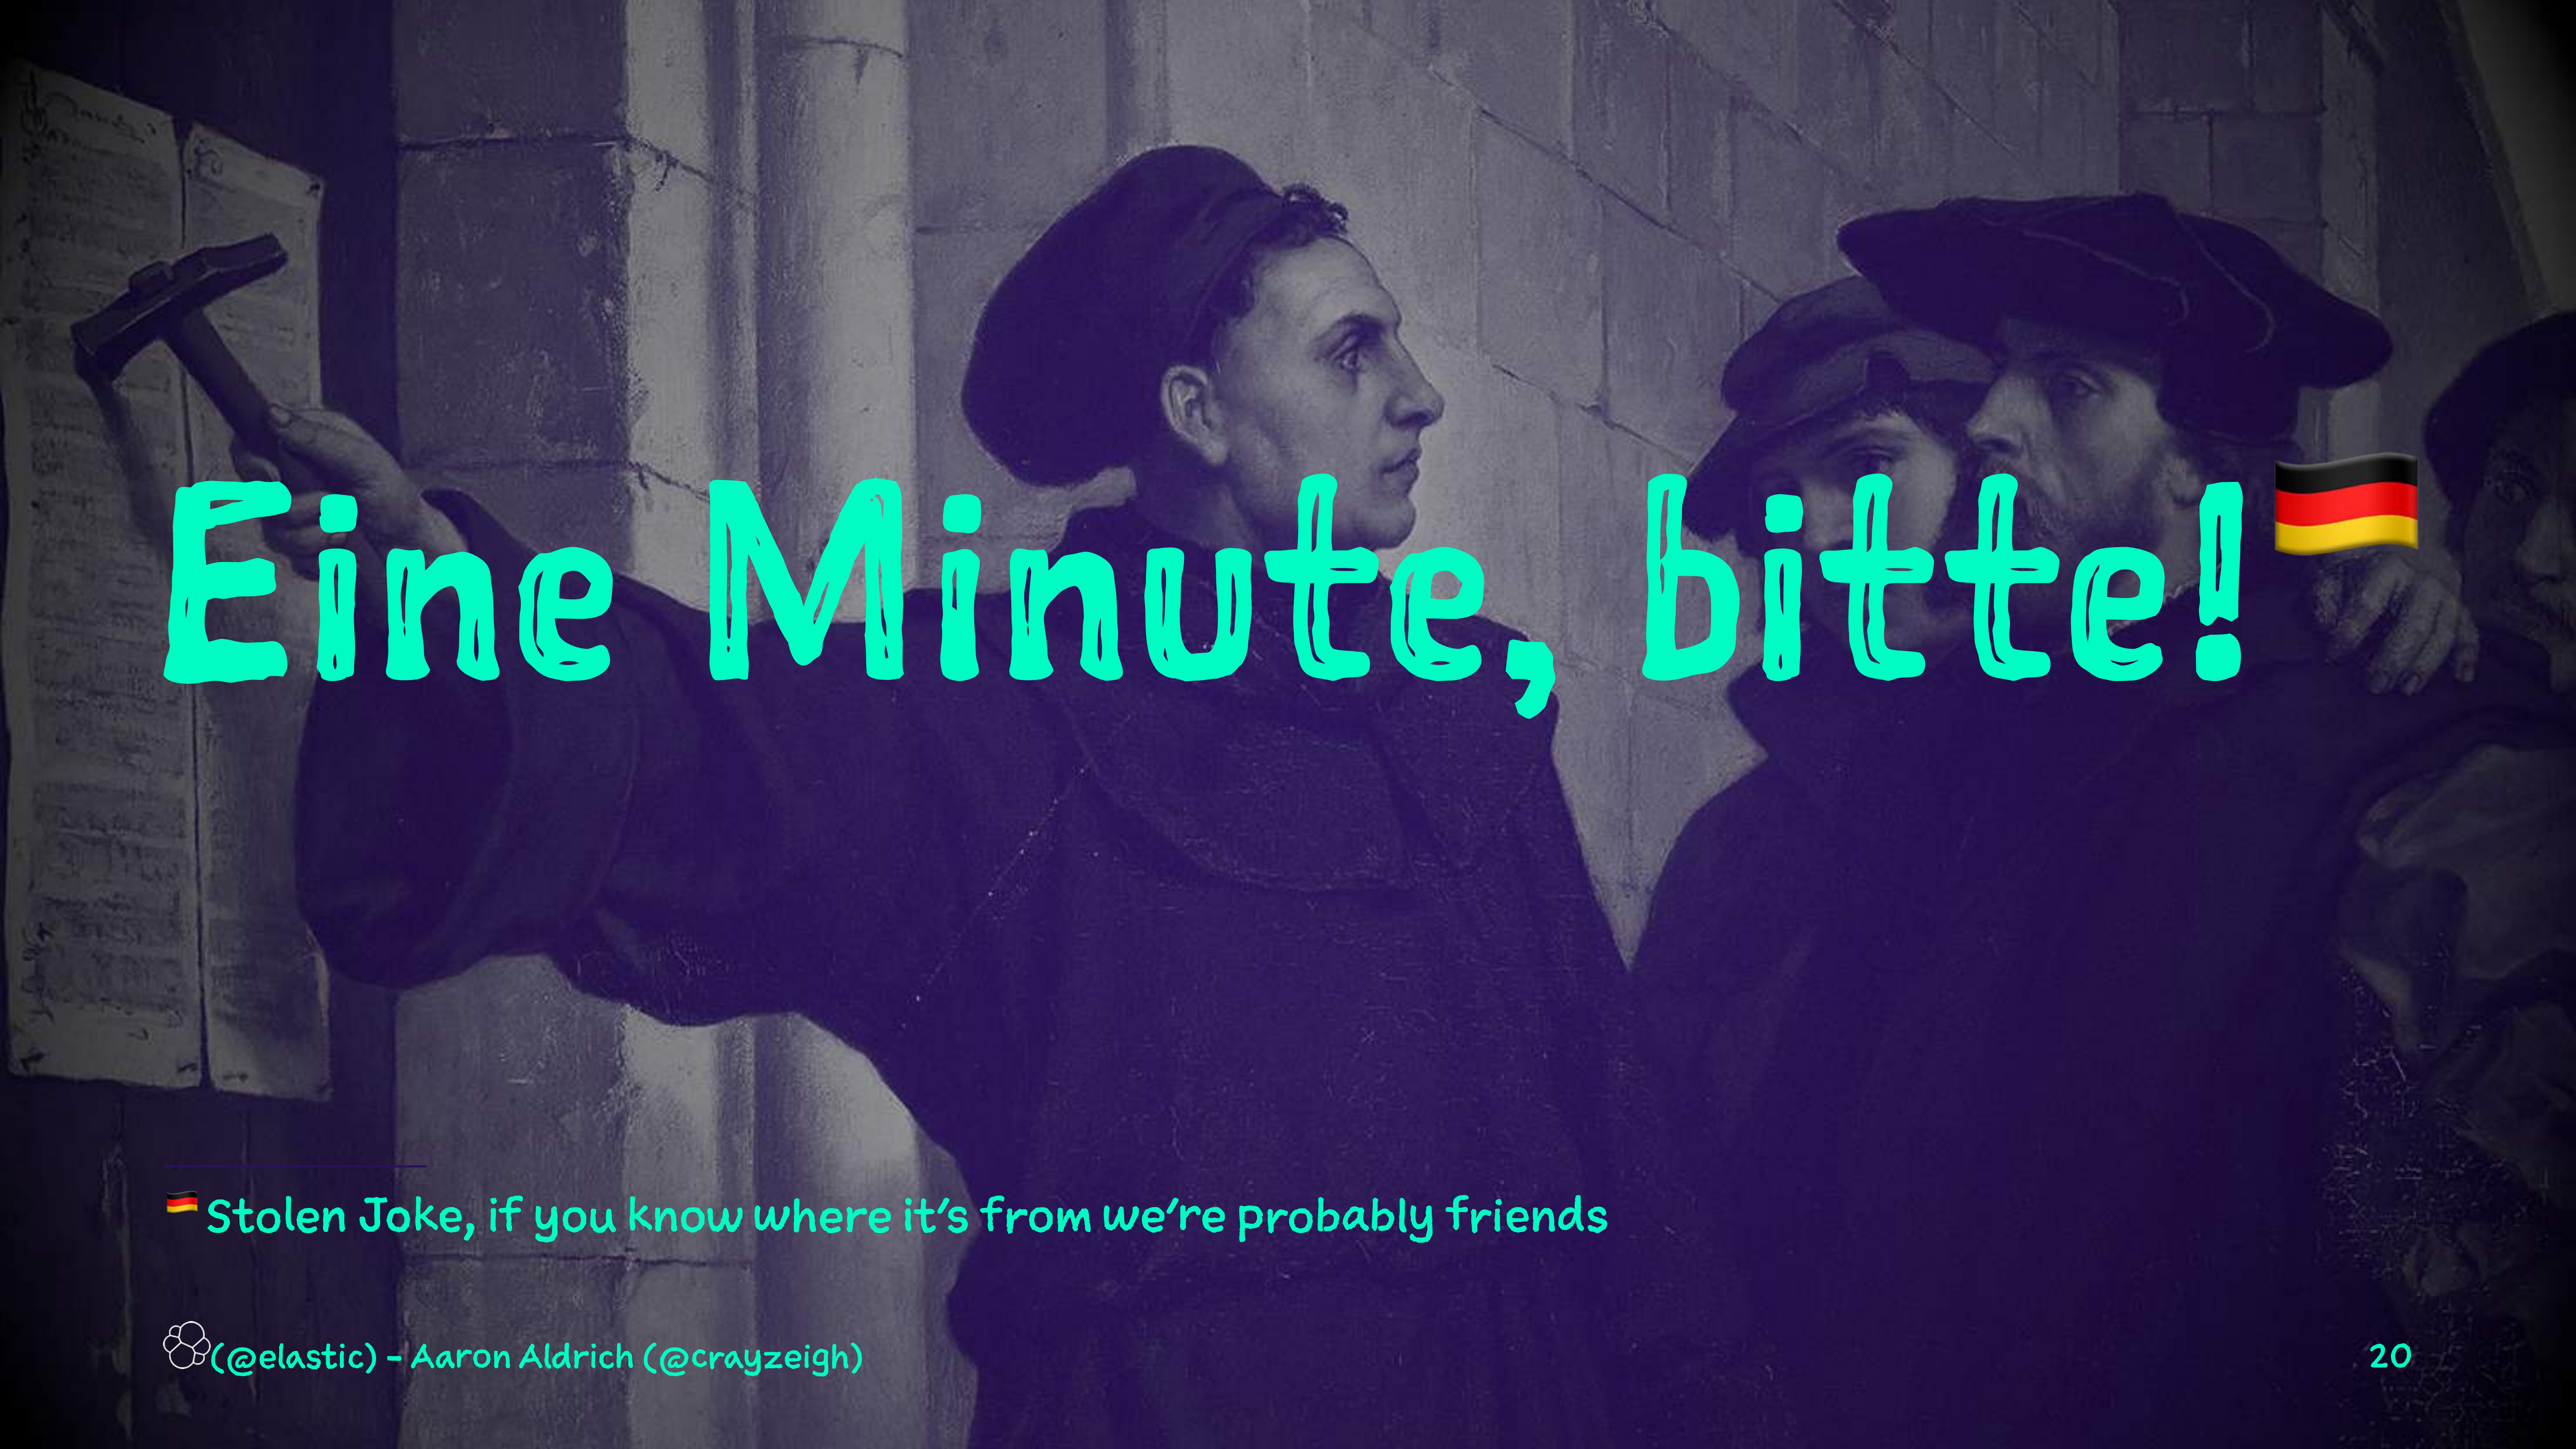

Eine Minute, bitte![🇩🇪]

[🇩🇪] Stolen Joke, if you know where it's from we're probably friends

NINES

don't matter...

NINES don't matter when USERS aren't HAPPY

~ Charity Majors (@mipsytipsy)

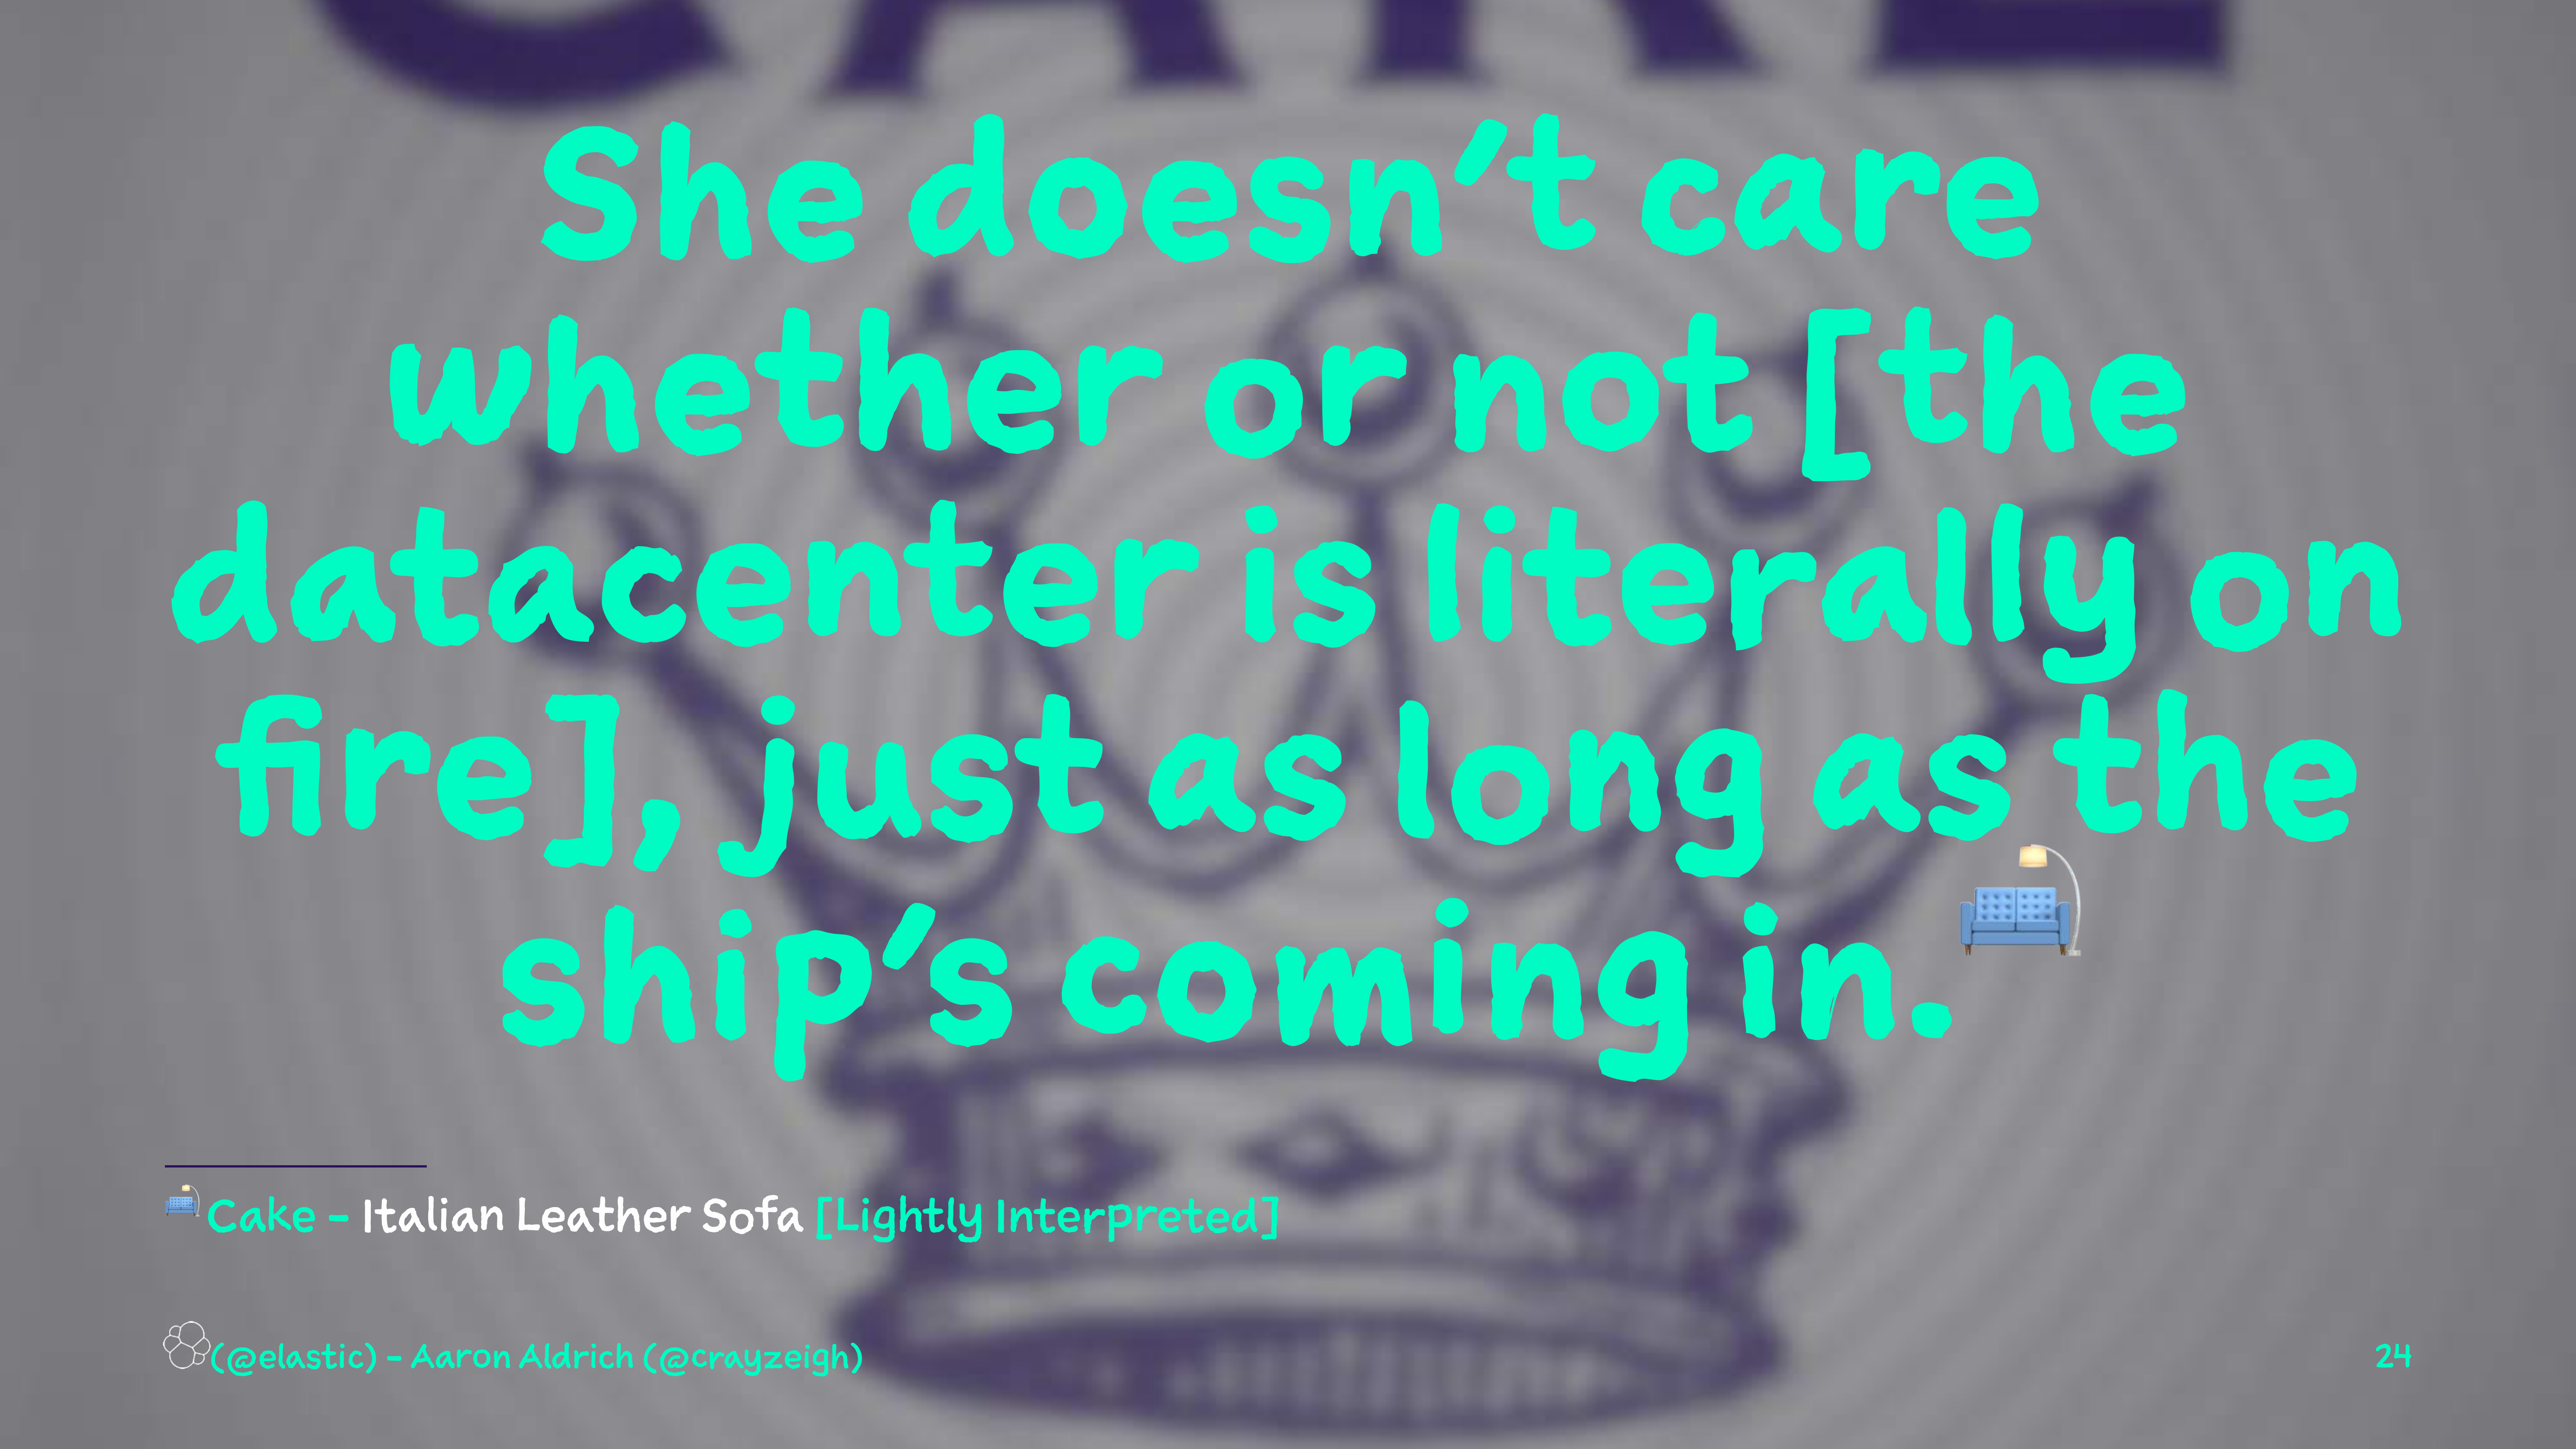

She doesn’t care whether or not [the datacenter is literally on fire], just as long as the ship’s coming in.🛋️

🛋️Cake - Italian Leather Sofa [Lightly Interpreted]

How does your business make money?

How do you help?

DevOps is about delivering Value

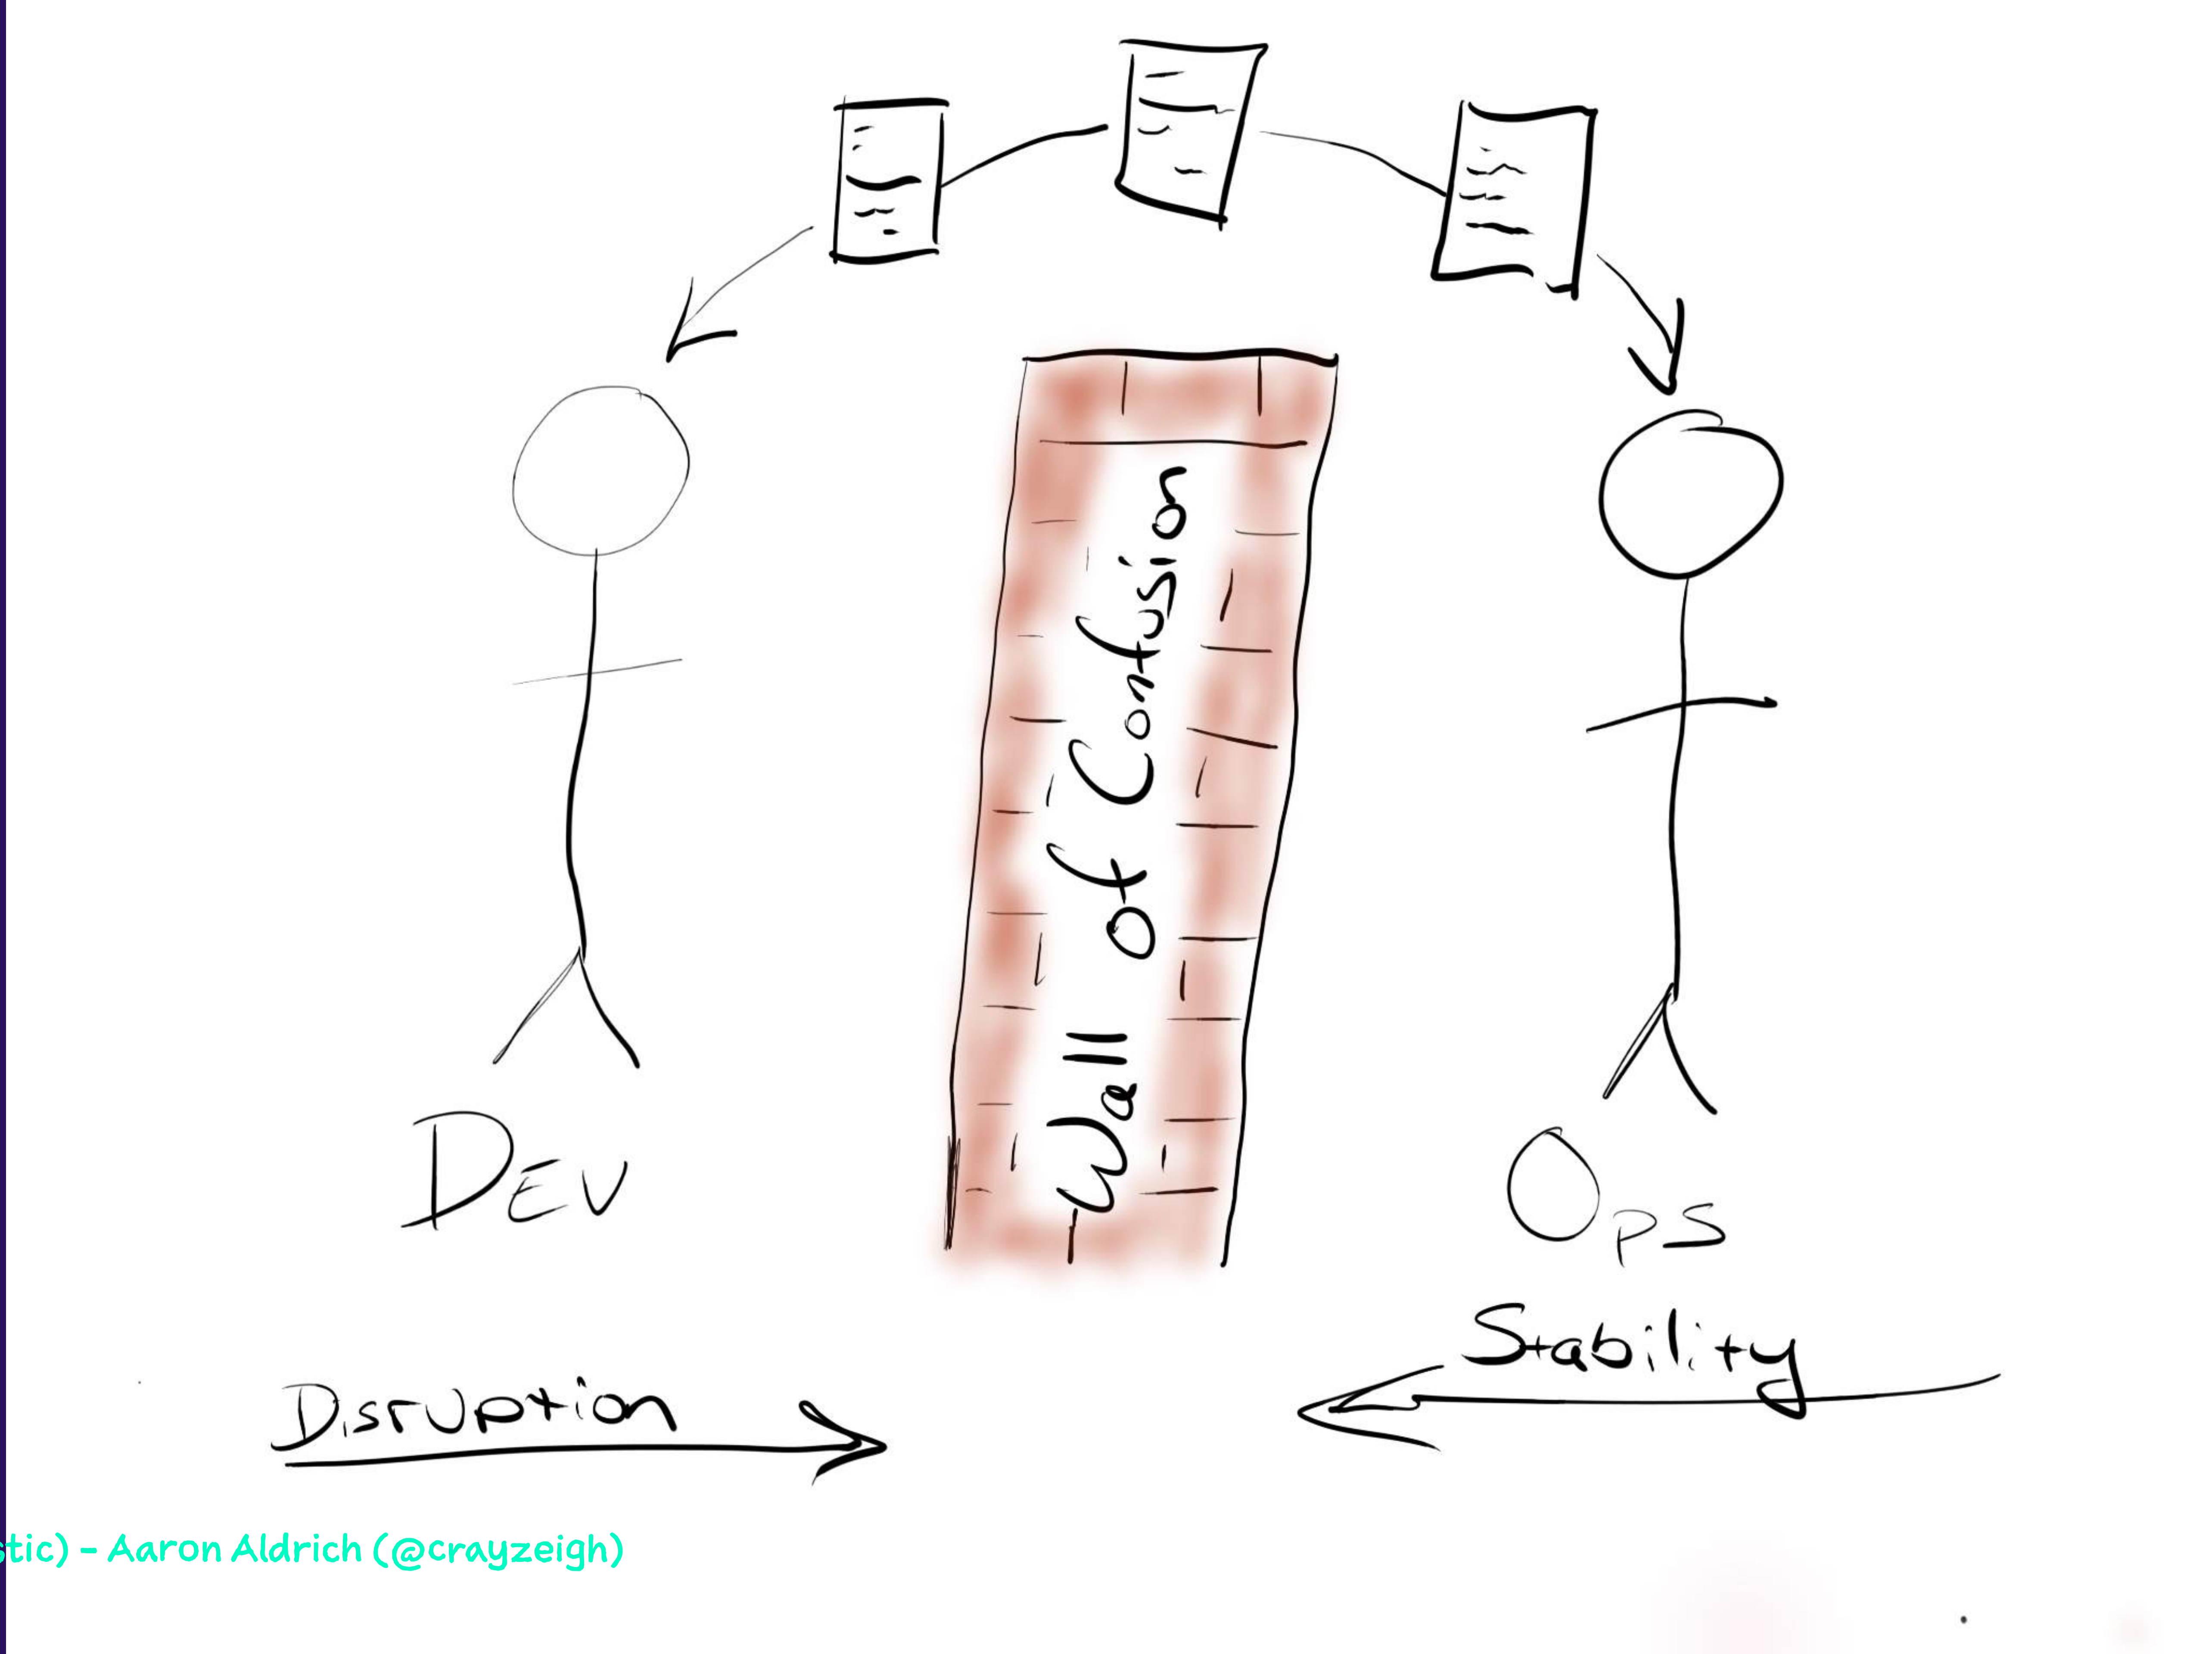
Dev
Wall of Confusion
Ops
Disruption
Stability

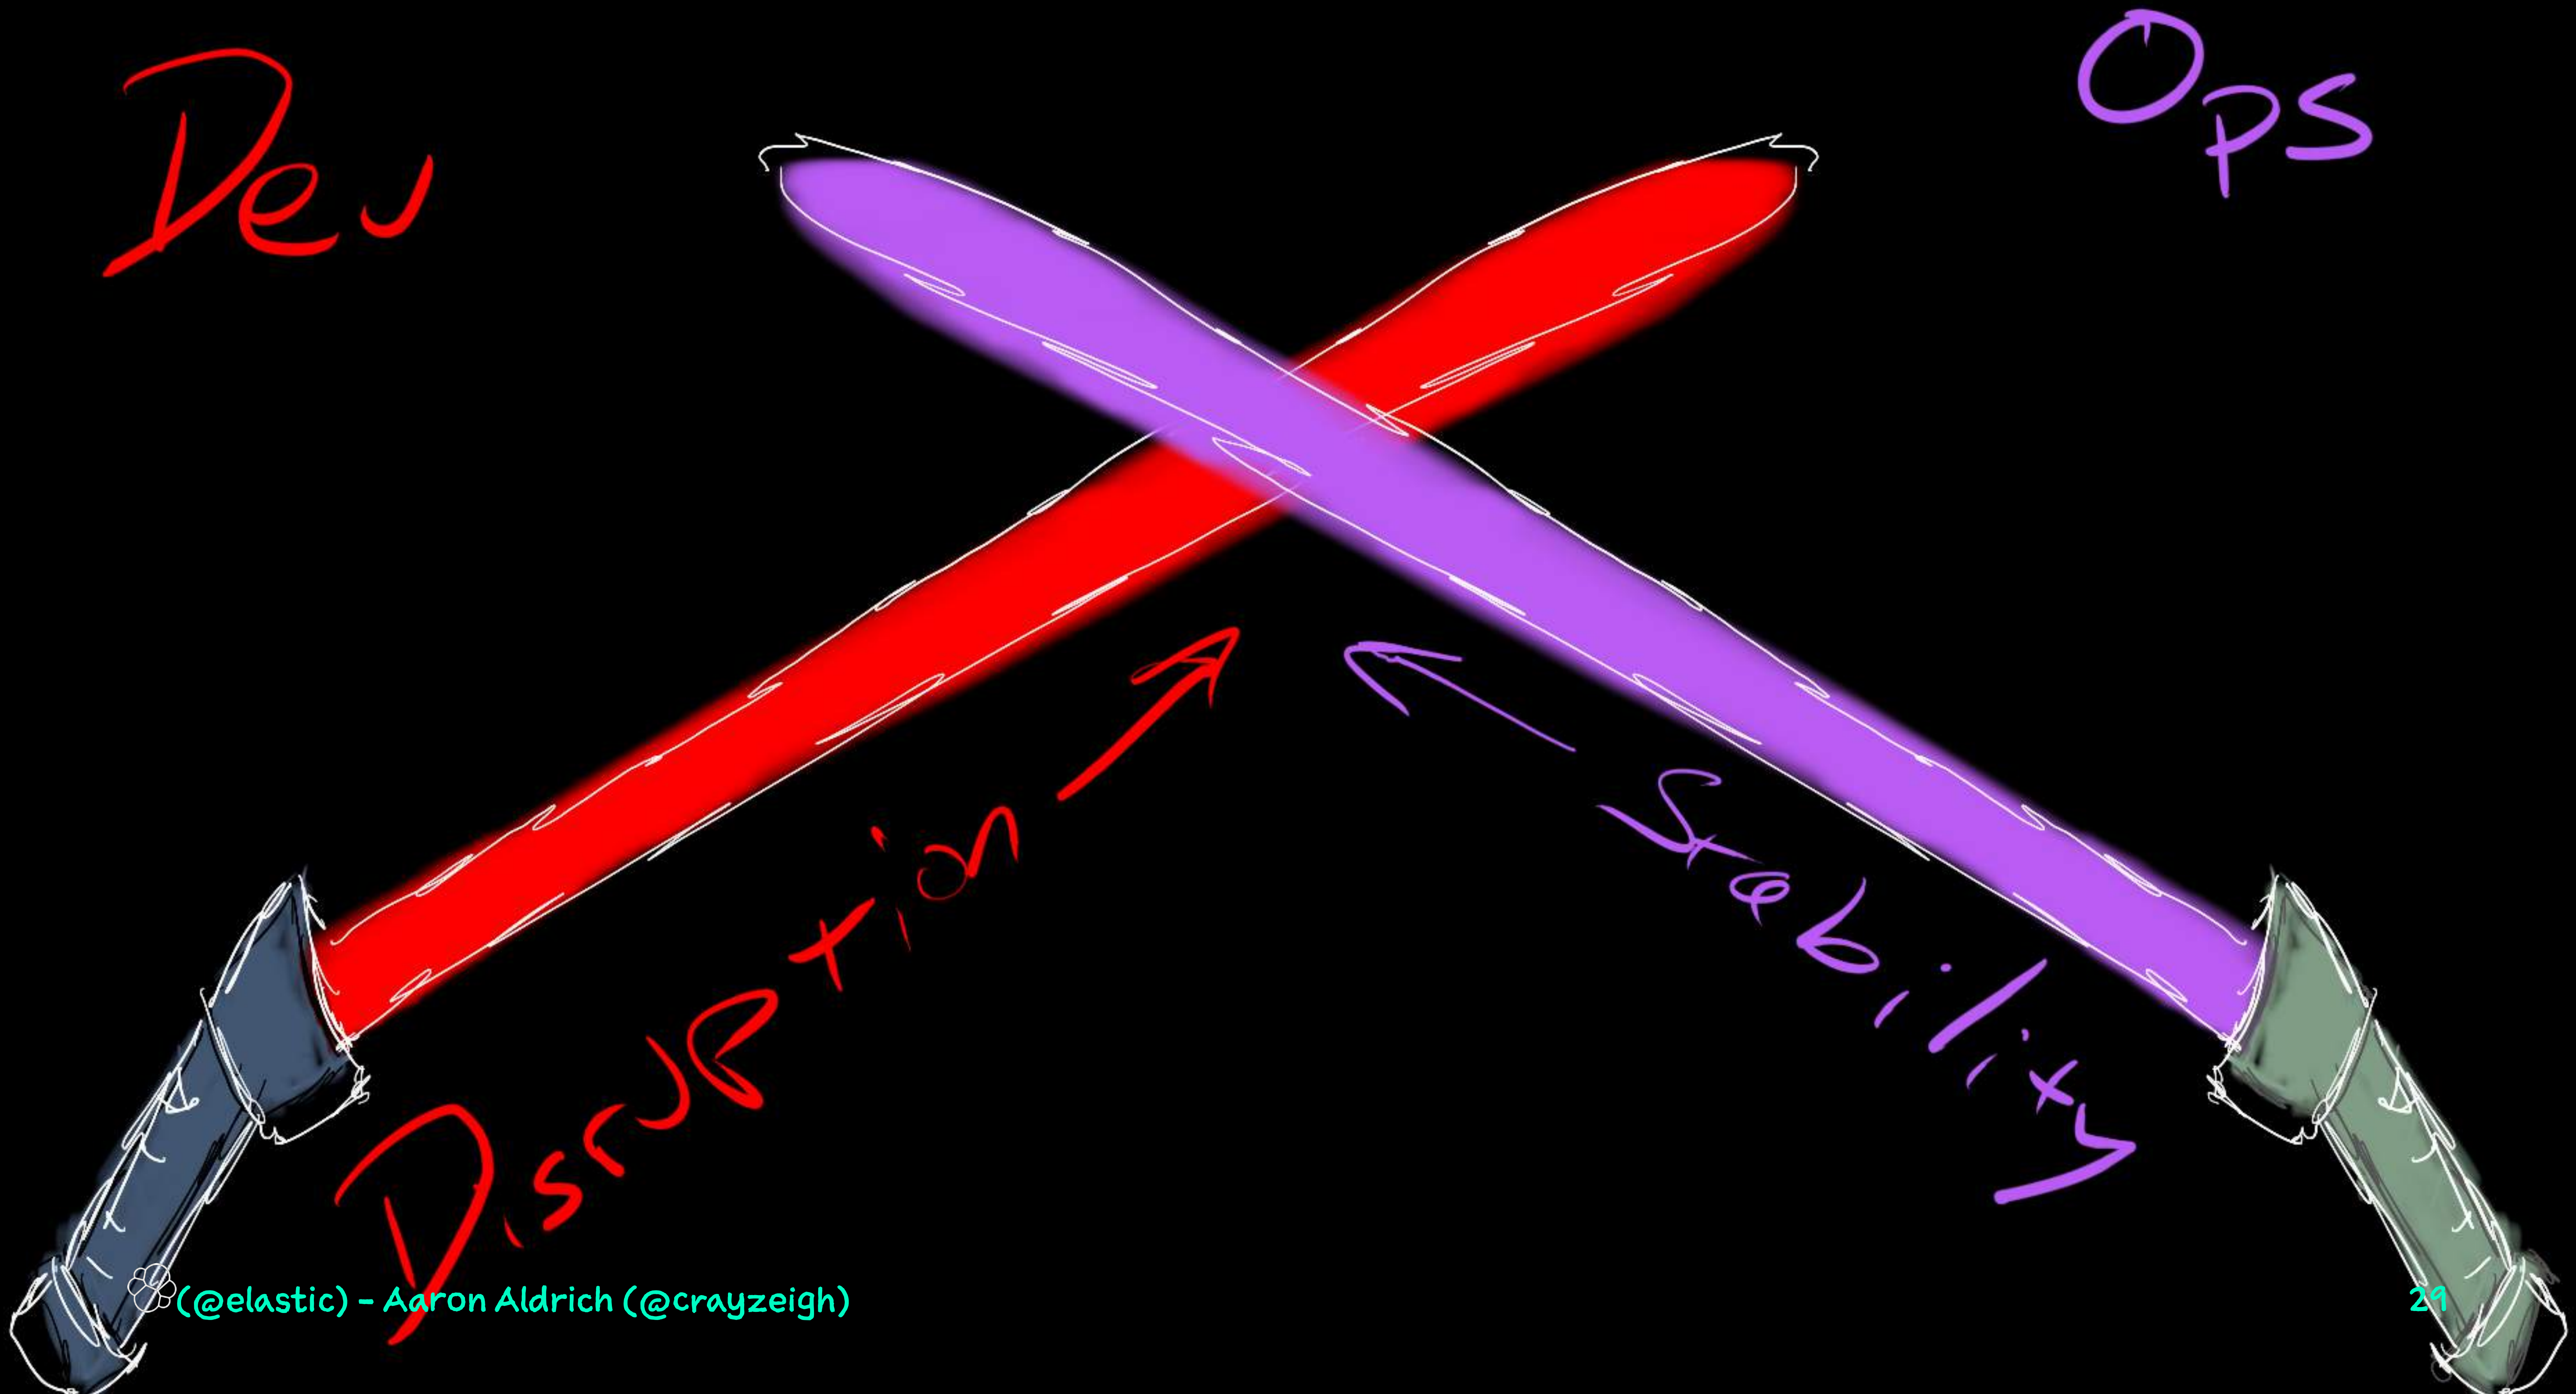
Dev
Ops
Disruption
Stability

Observability

Isn't it just monitoring with better SEO?
- You

You're not wrong...

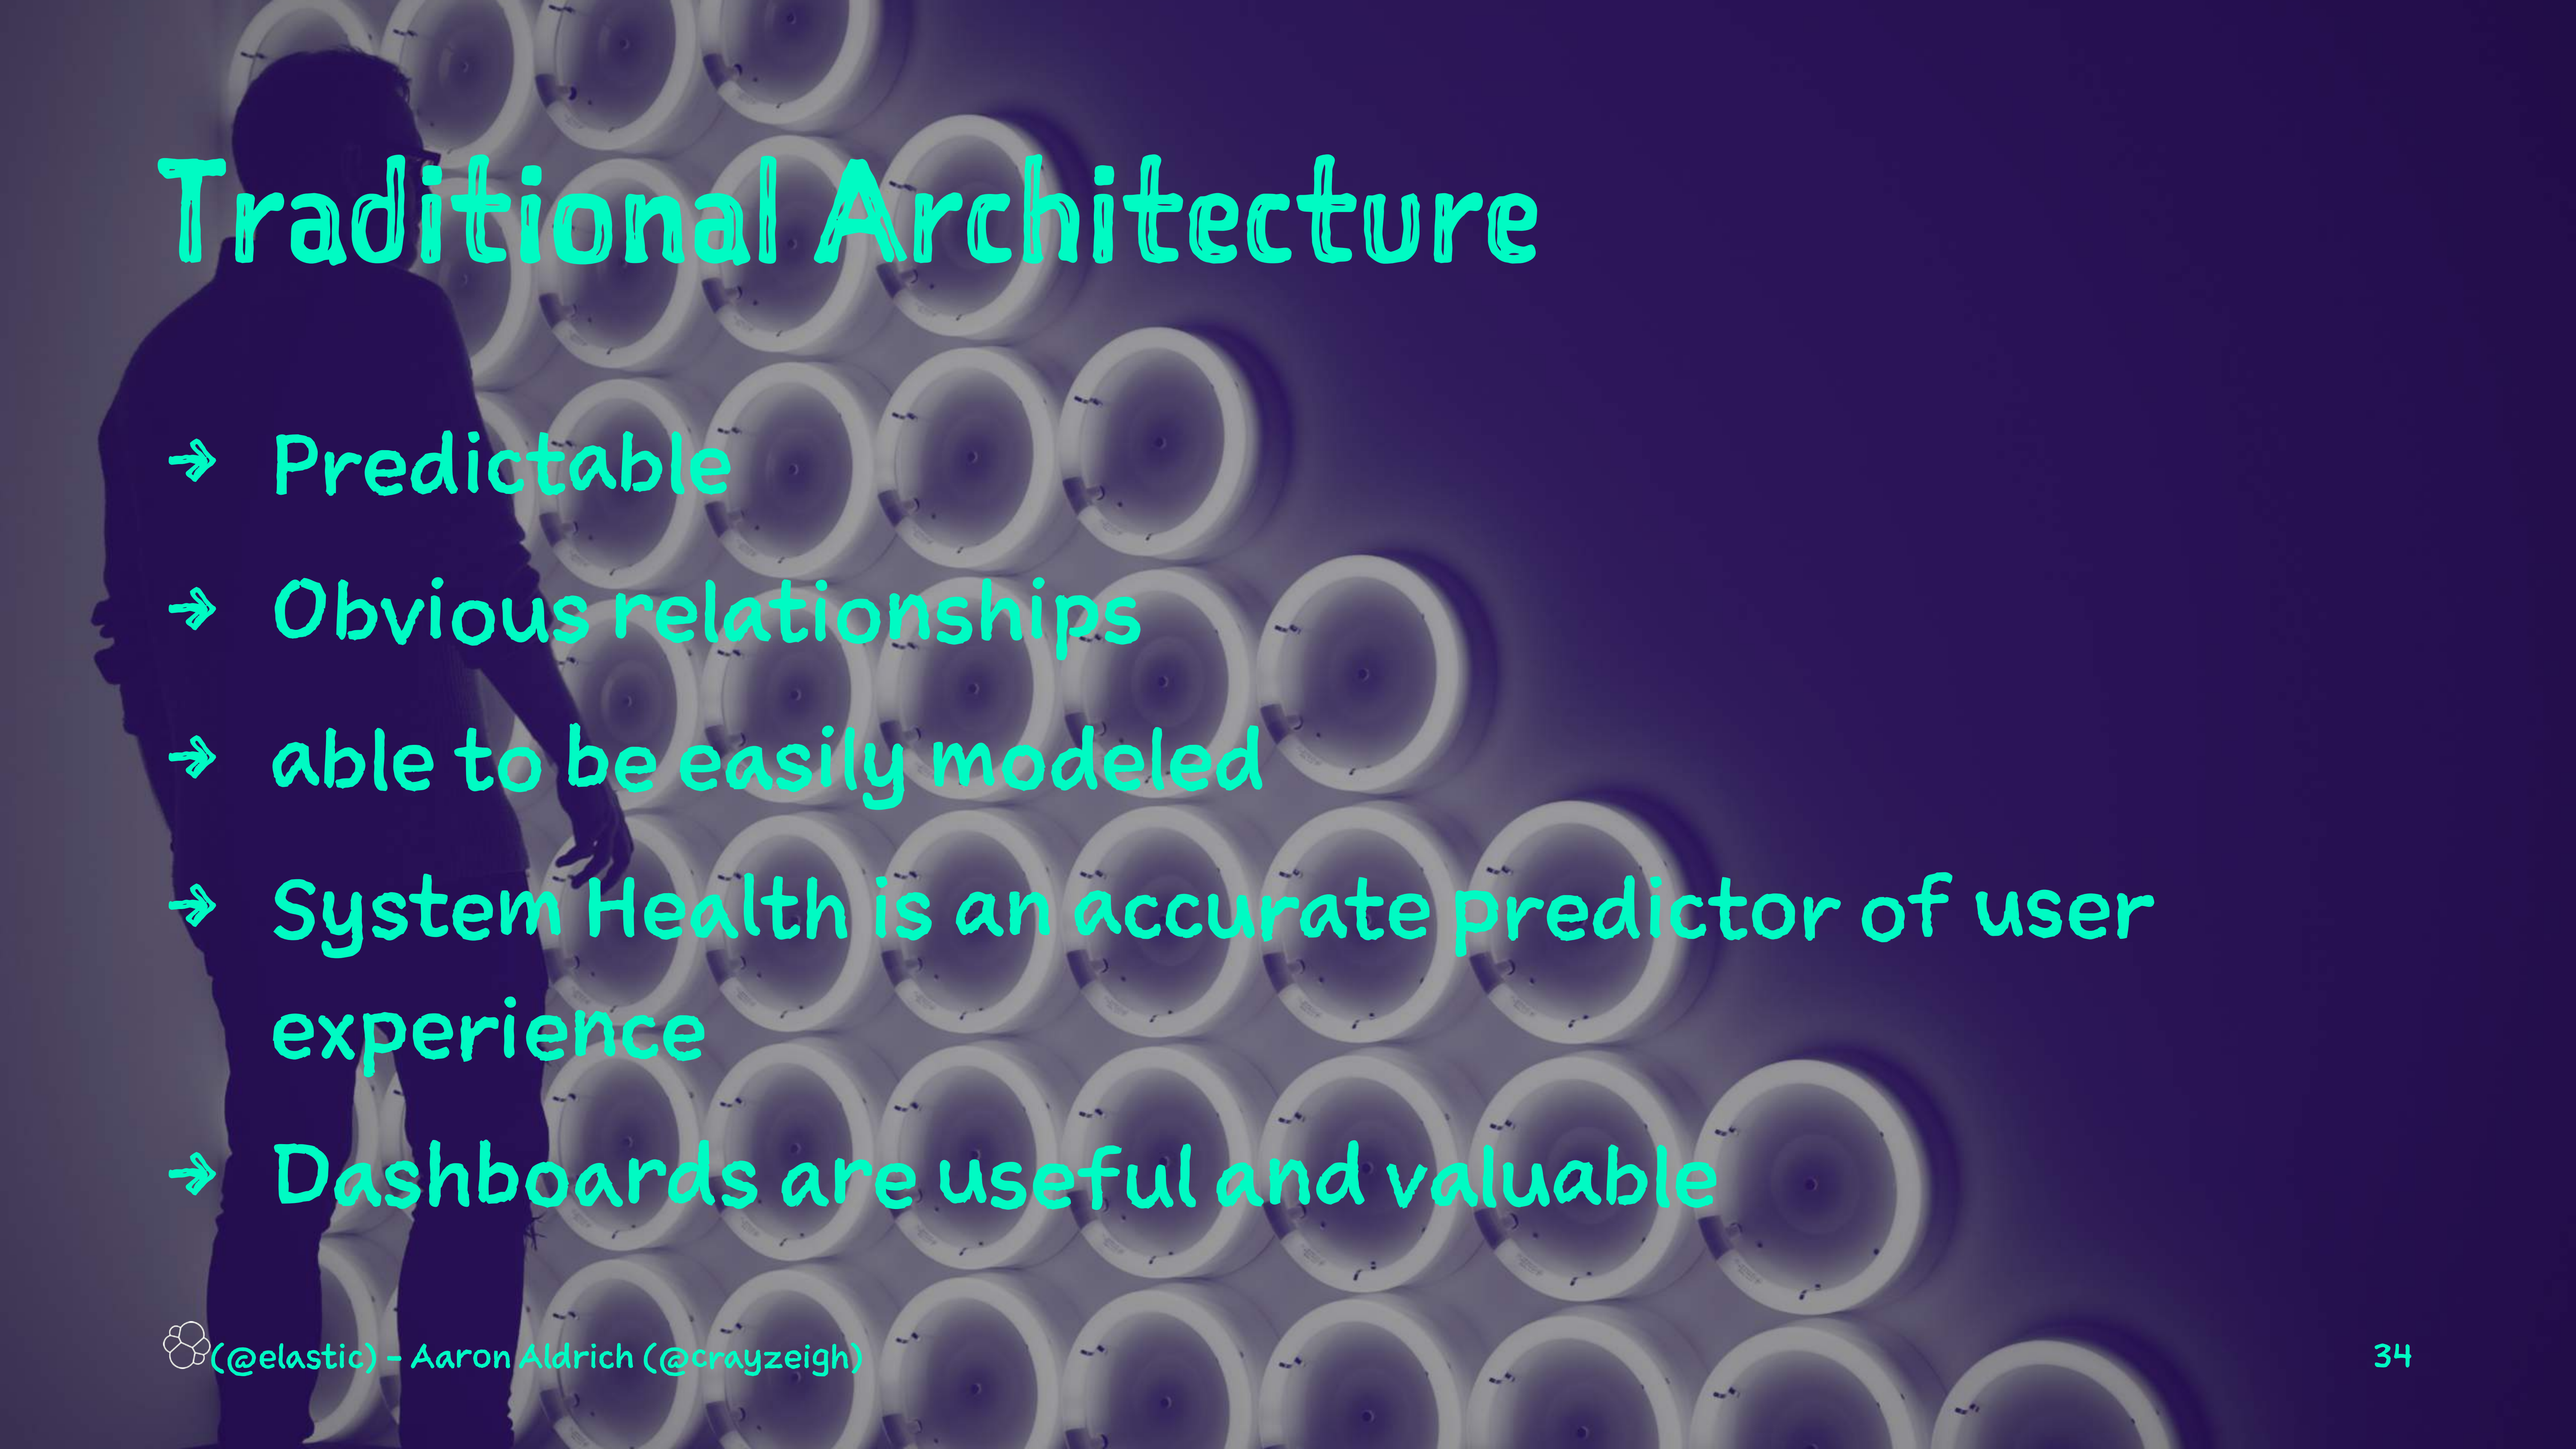

Traditional Architecture

- Predictable
- Obvious relationships
- able to be easily modeled
- System Health is an accurate predictor of user experience
- Dashboards are useful and valuable

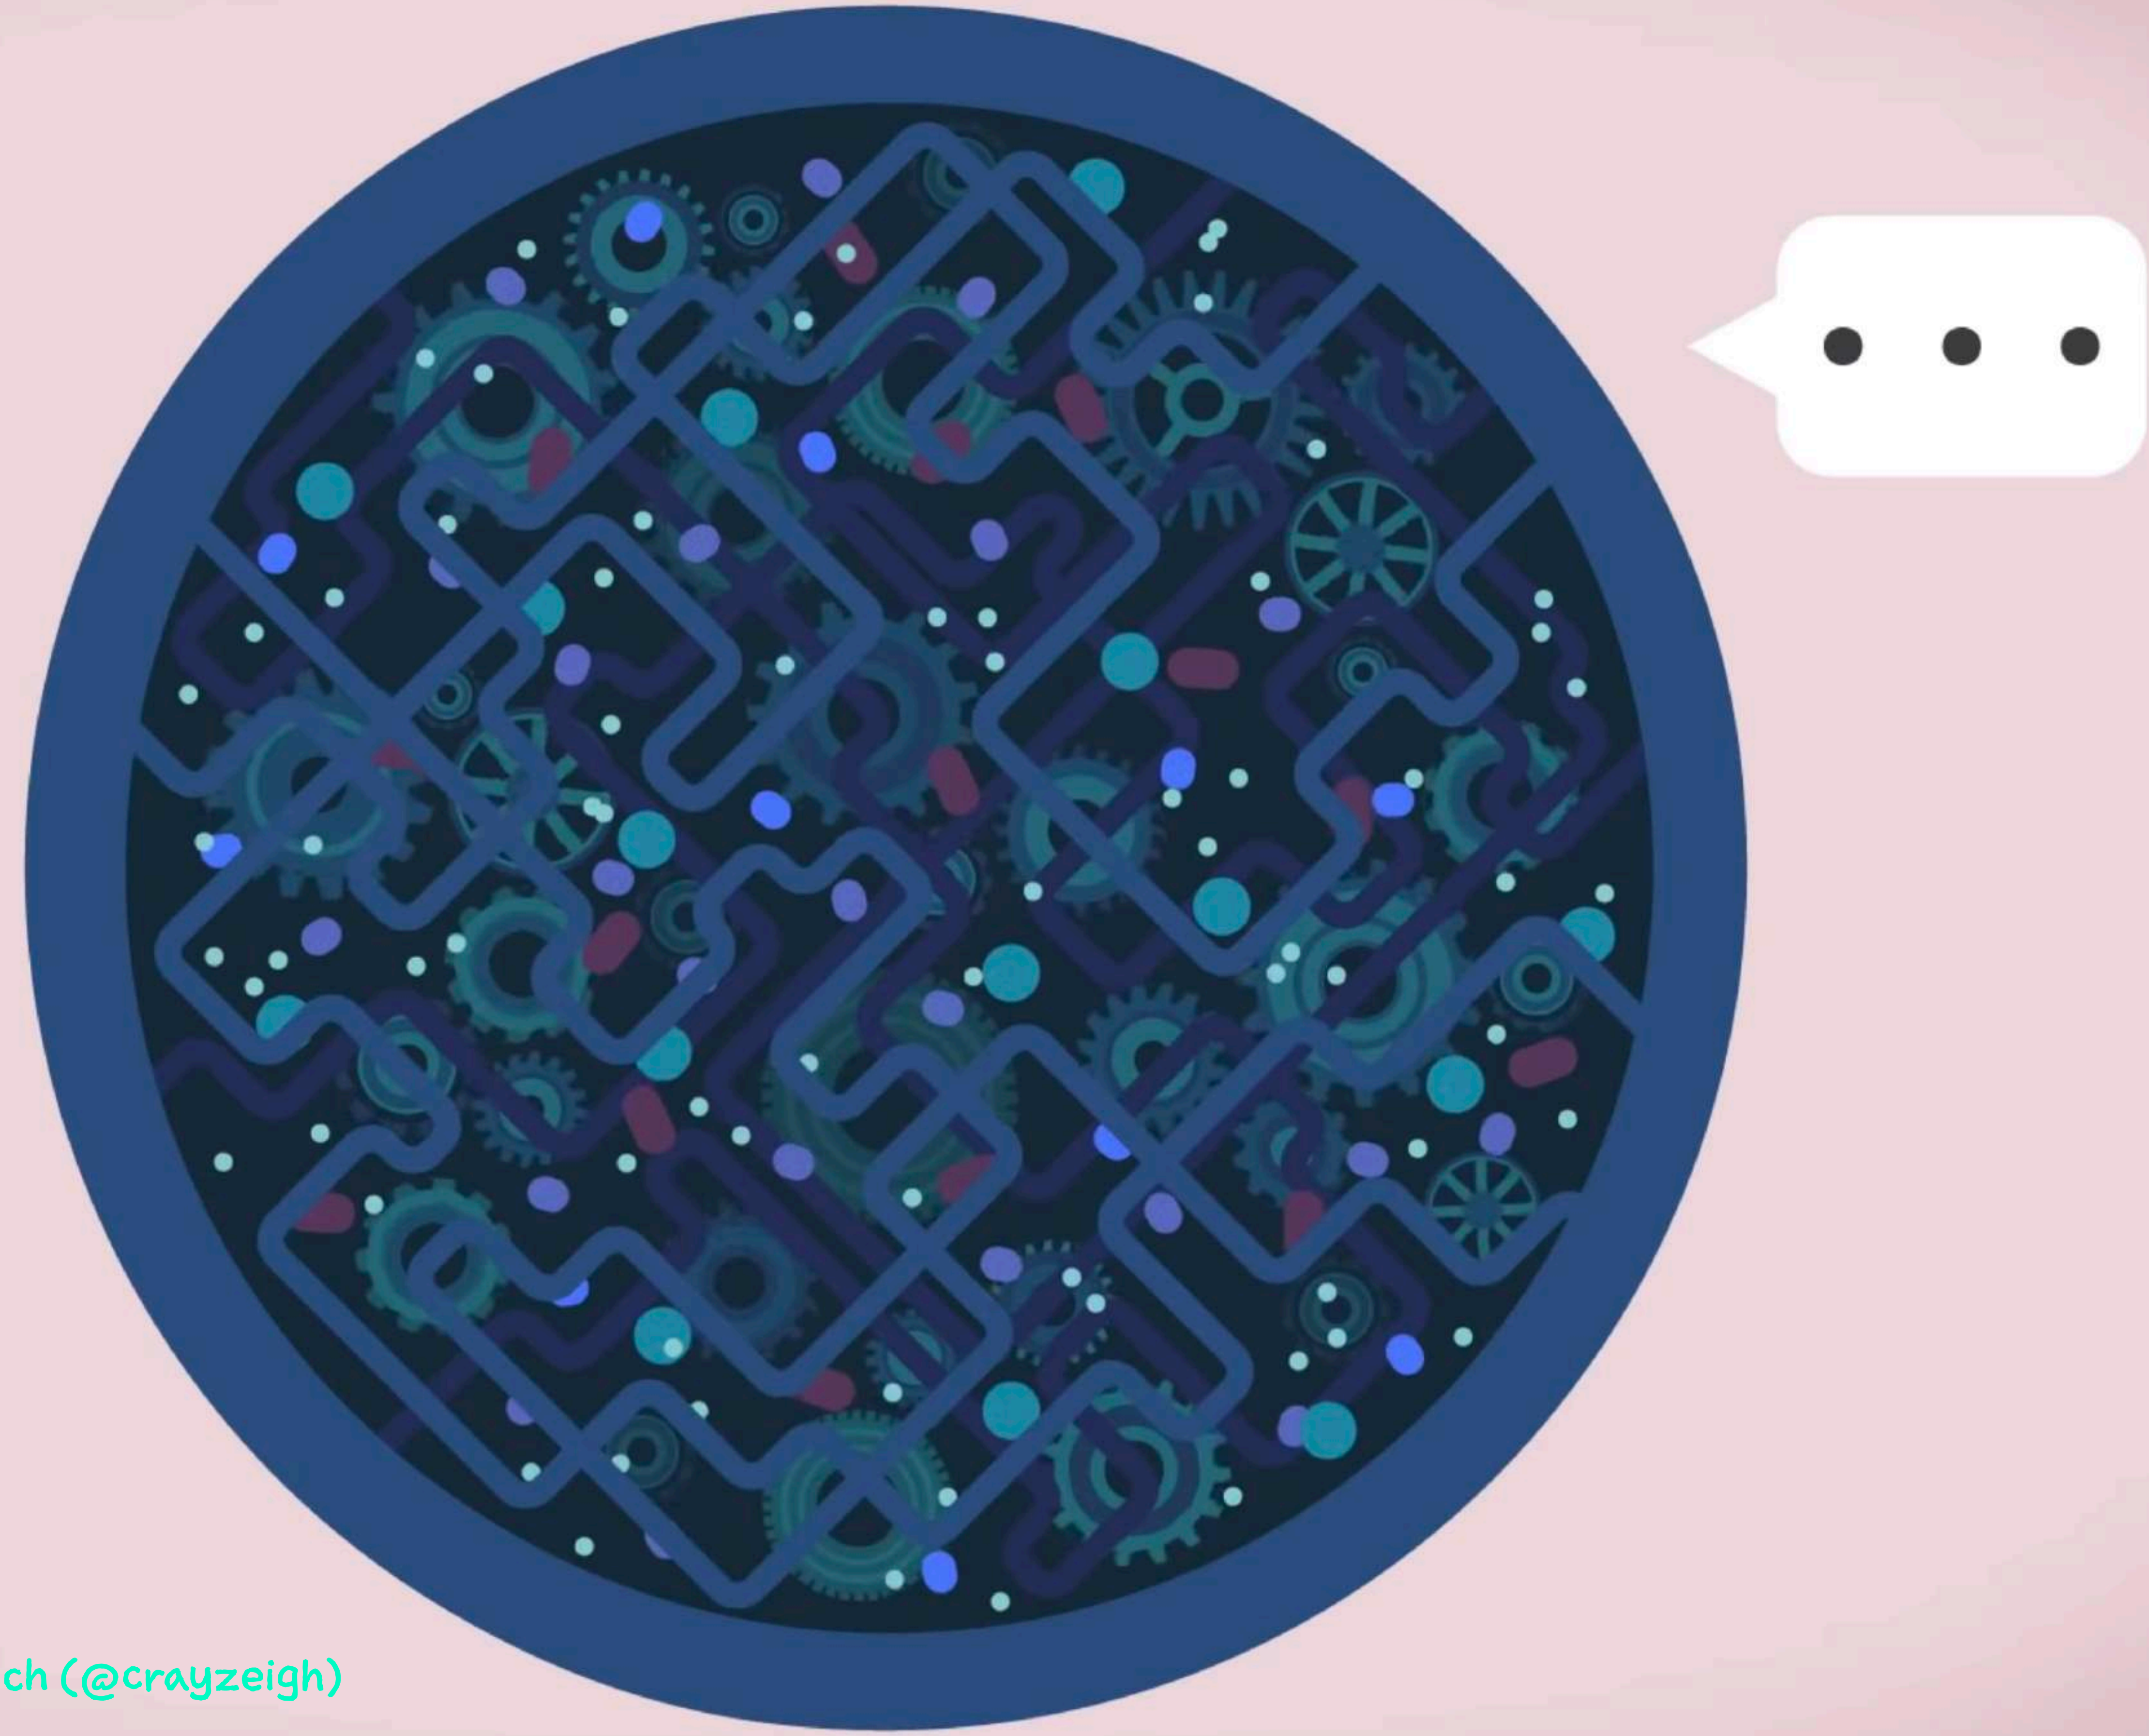

Complex Systems

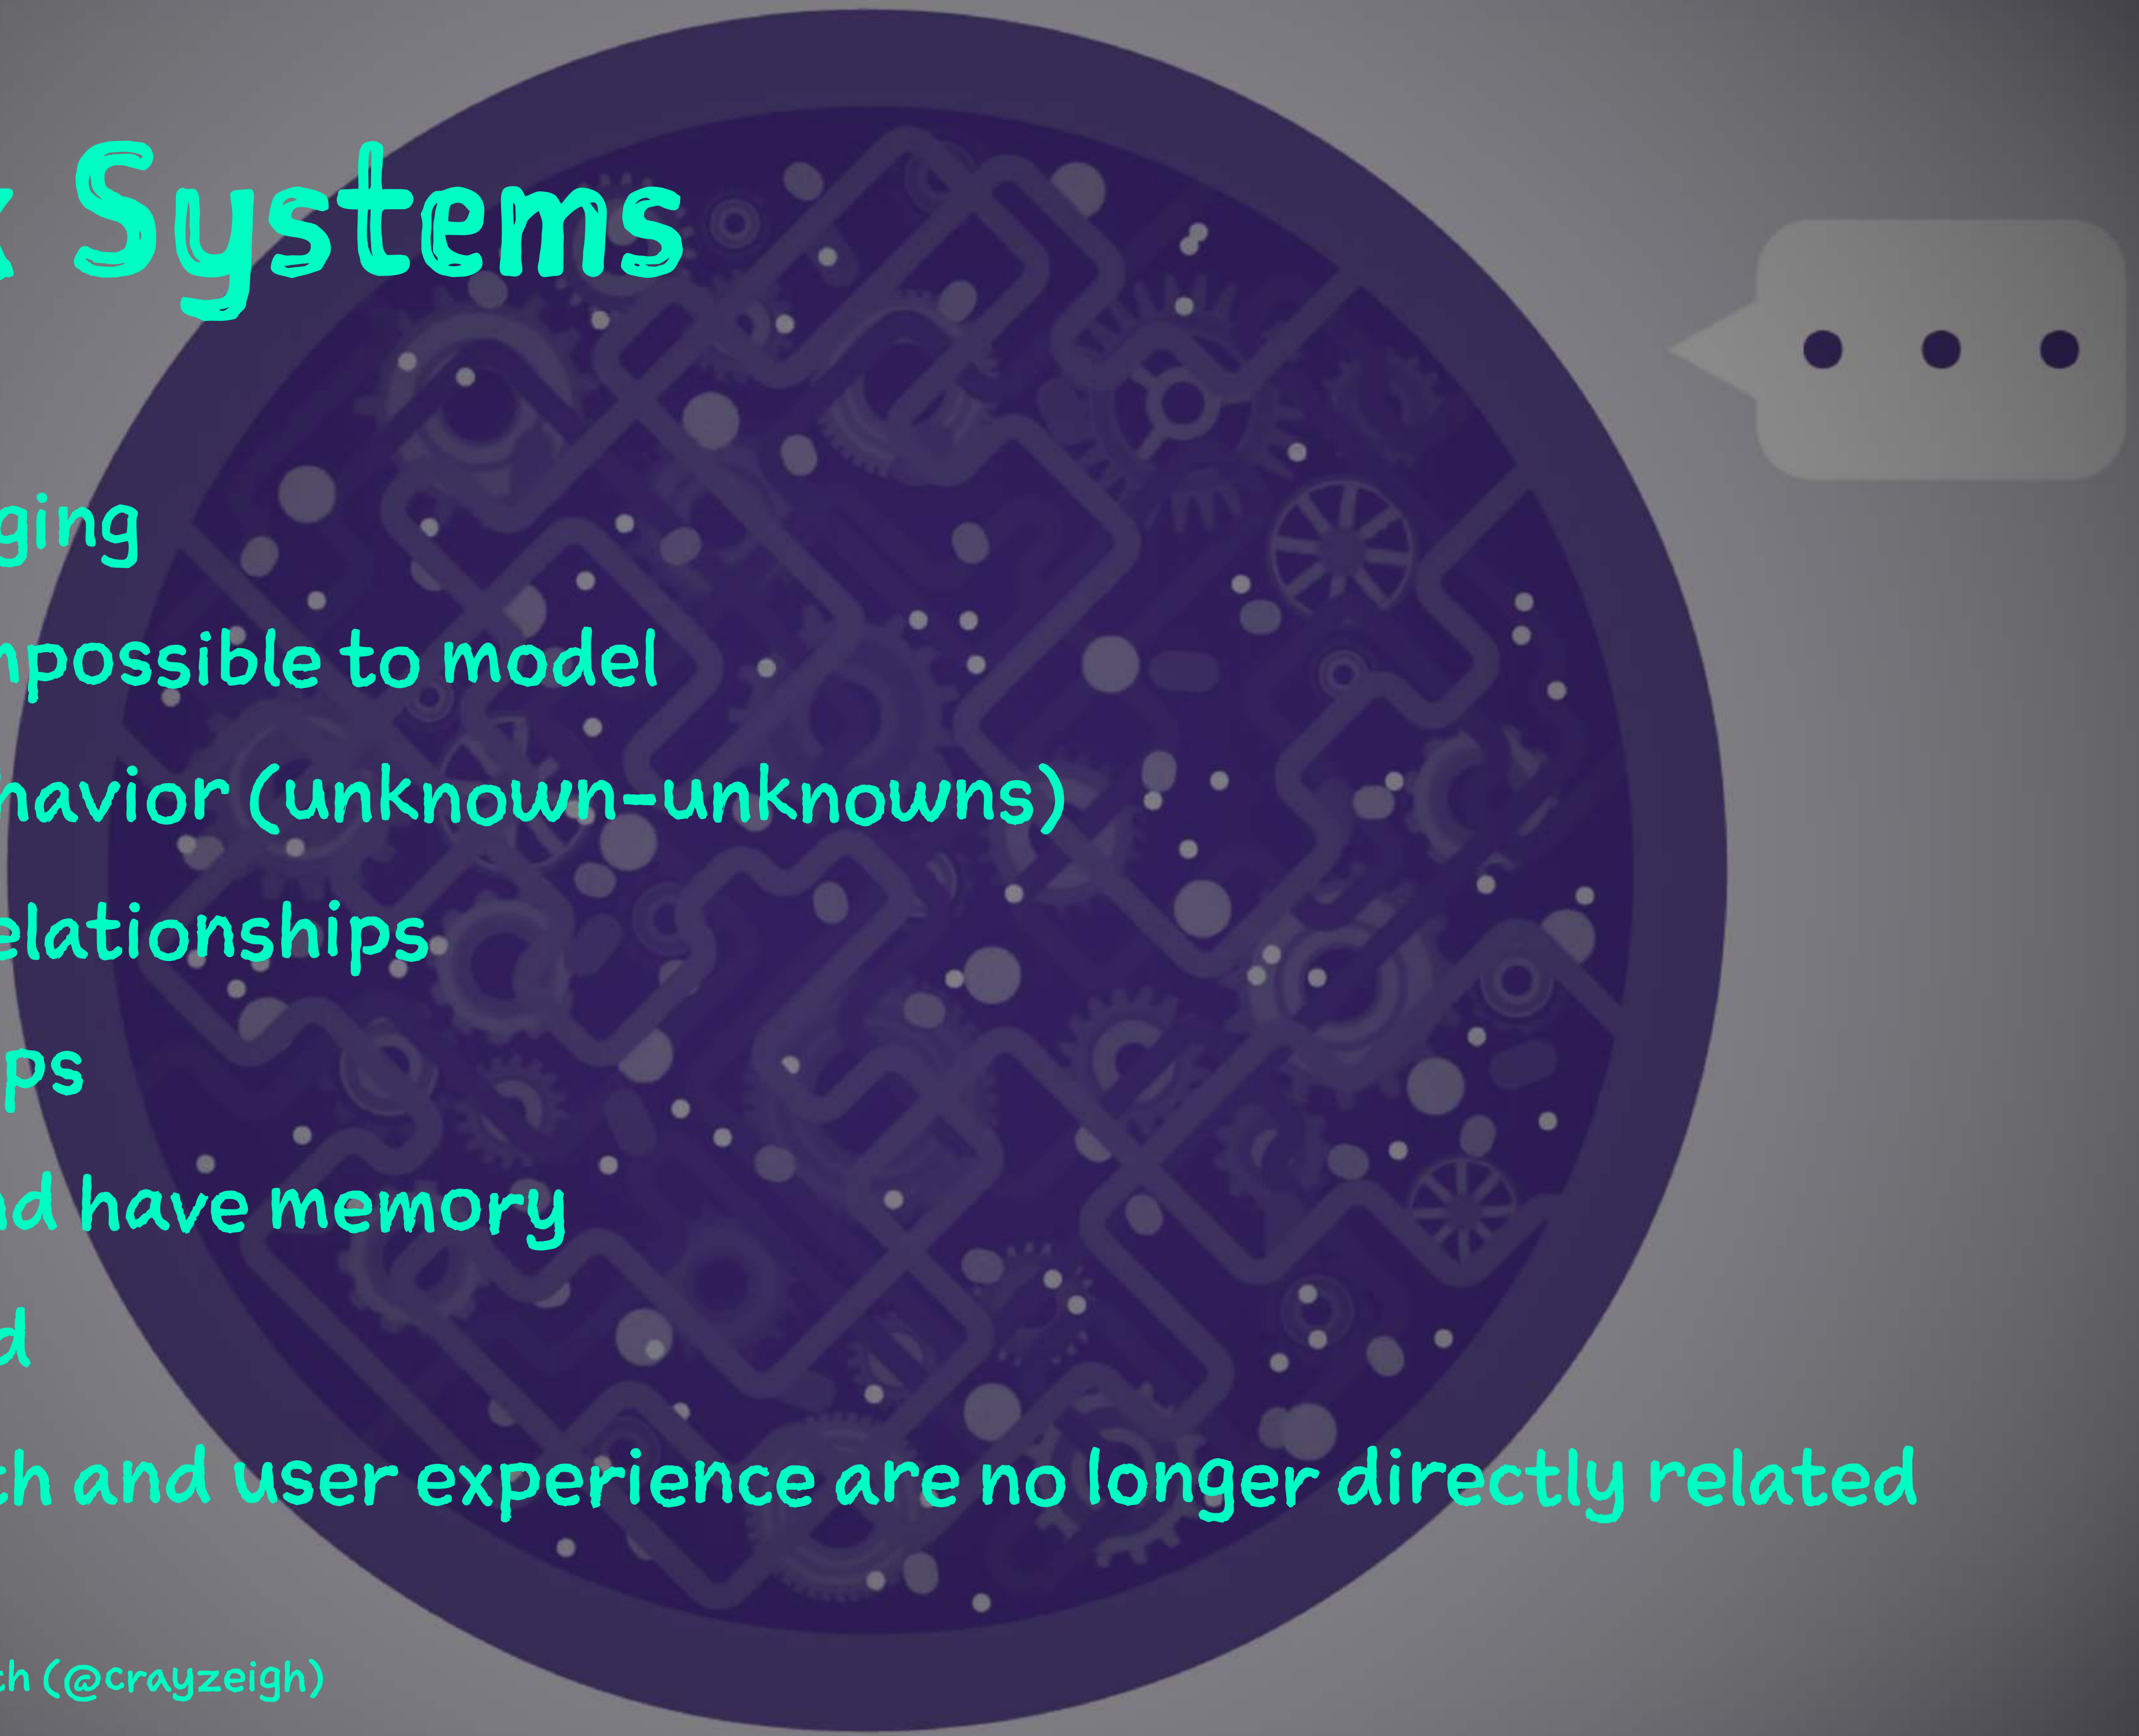

- Always changing
- Difficult or impossible to model
- emergent behavior (unknown-unknowns)
- non-linear relationships
- feedback loops
- can adapt and have memory
- can be nested
- System health and user experience are no longer directly related

Root Cause

is a myth

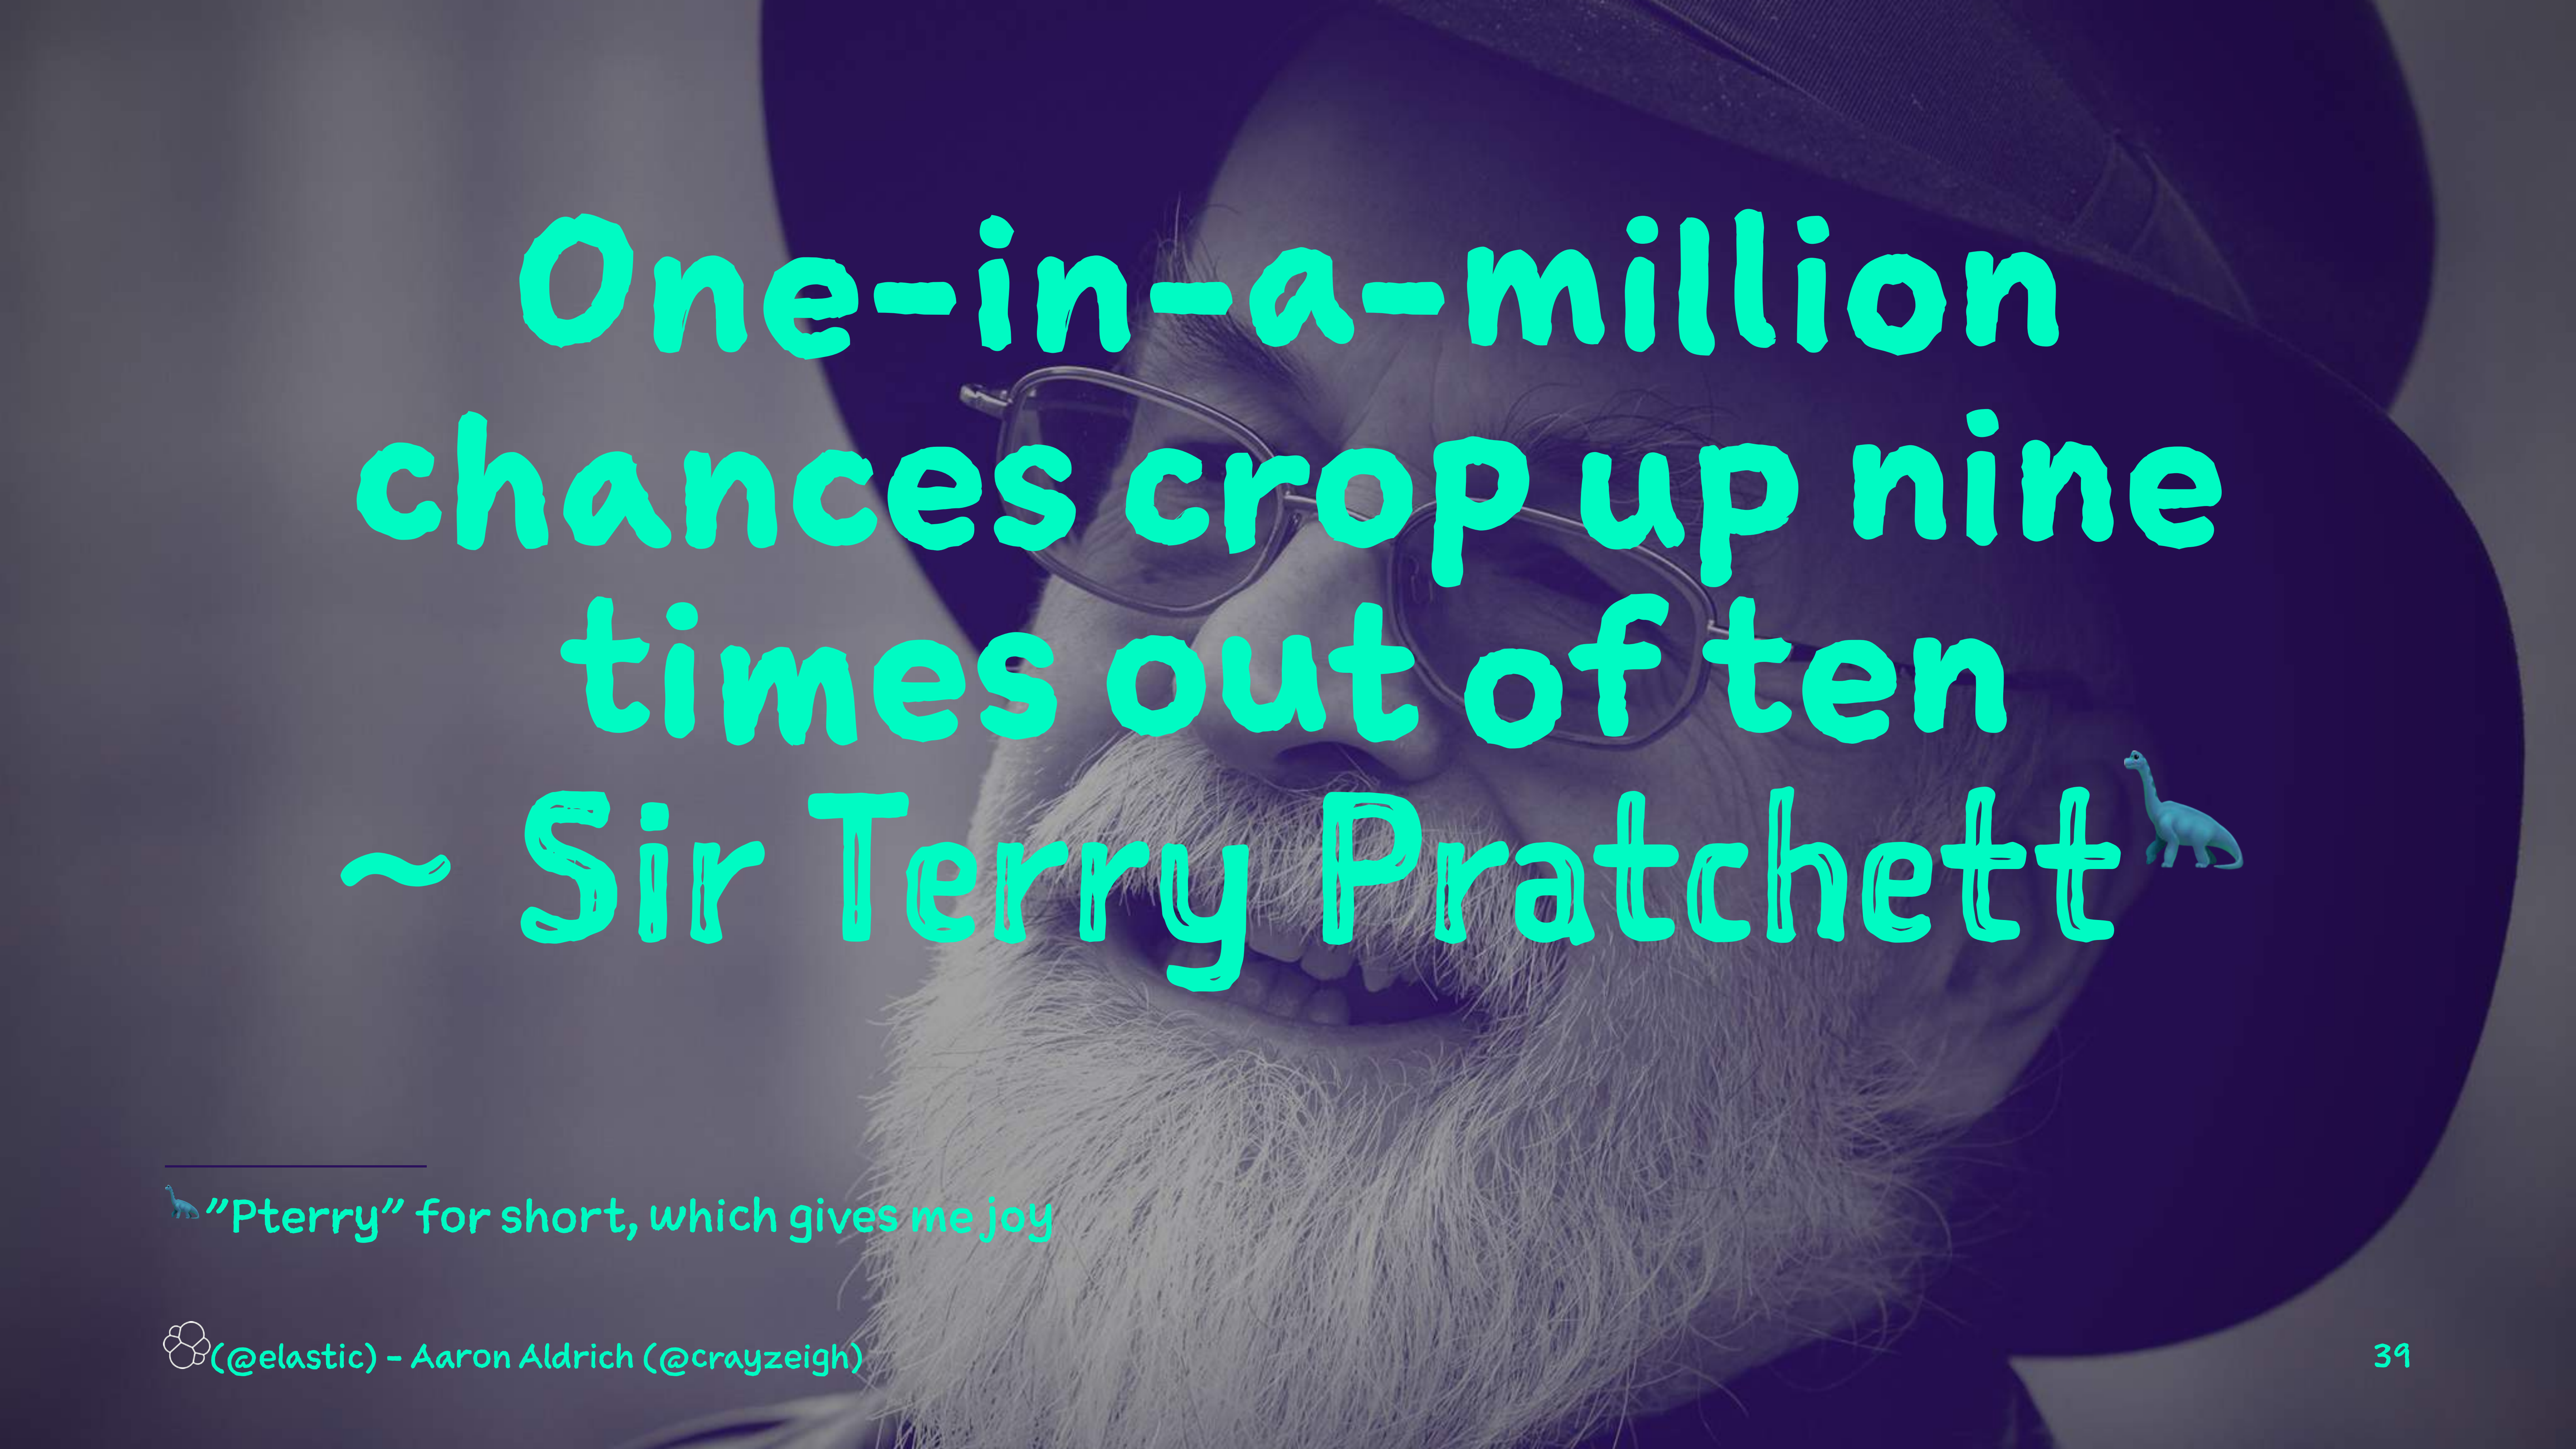

One-in-a-million chances crop up nine times out of ten
~ Sir Terry Pratchett🦕

🦕"Pterry" for short, which gives me joy

SRE

SLI
SLO
SLA

Services not systems

Site Reliability Engineering

- (SLI) What is availability?
- (SLO) How much do we actually need?
- (SLA) What happens when we're not meeting this target?

Site Reliability Engineering

- (SLI) What is availability?
- (SLO) How much do we actually need?
- ~~(SLA) What happens when we're not meeting this target?~~

Service Level Indicators

- Is it up?
- 200OK
- latency
- percentiles or medians for meaning

Service Level Indicators

- Is it up?
- 200OK
- latency
- percentiles or medians[😏] for meaning

[😏] Never trust averages, they hide data

Never trust averages, they hide data

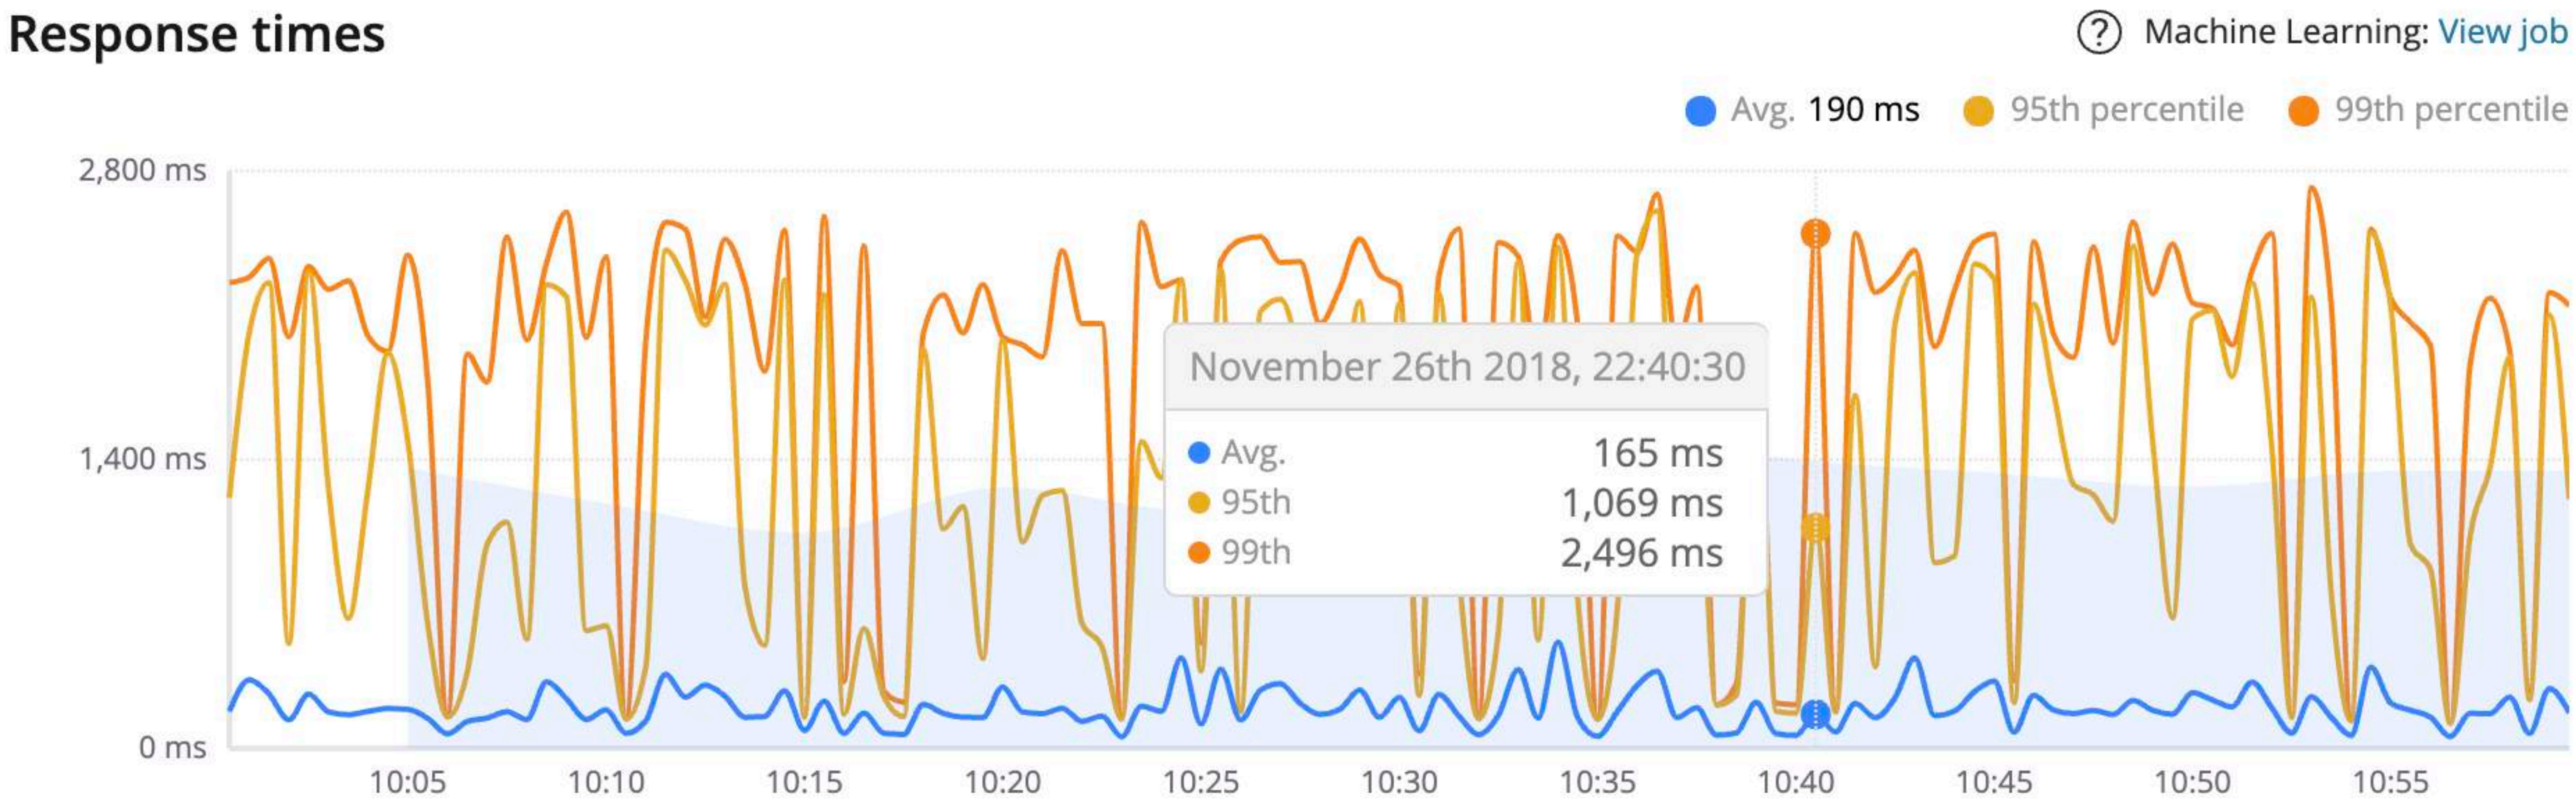

November 26th 2018, 22:40:30

● Avg.	165 ms
● 95th	1,069 ms
● 99th	2,496 ms

The 99th percentile latency of requests received in <300 ms and responded to with a 200 status

Service Level Objectives

How much availability do we need?

99%

99.9%

99.99%

99.999%

Each 9 is exponentially more expensive to provide

availability	avg per year	avg per day
99%	3.65 days	14.4 minutes
99.9%	8.76 hours	1.44 minutes
99.99%	52.56 minutes	8.64 seconds
99.999%	5.25 minutes	863 ms

A good SLO barely keeps users happy

(these should be driving your alerts)

Error Budgets

It's GOOD to have errors

Error Budgets

Bring Balance to the Force

SLAs = 💵

~~SLAs~~ = 💸

What about the fire?

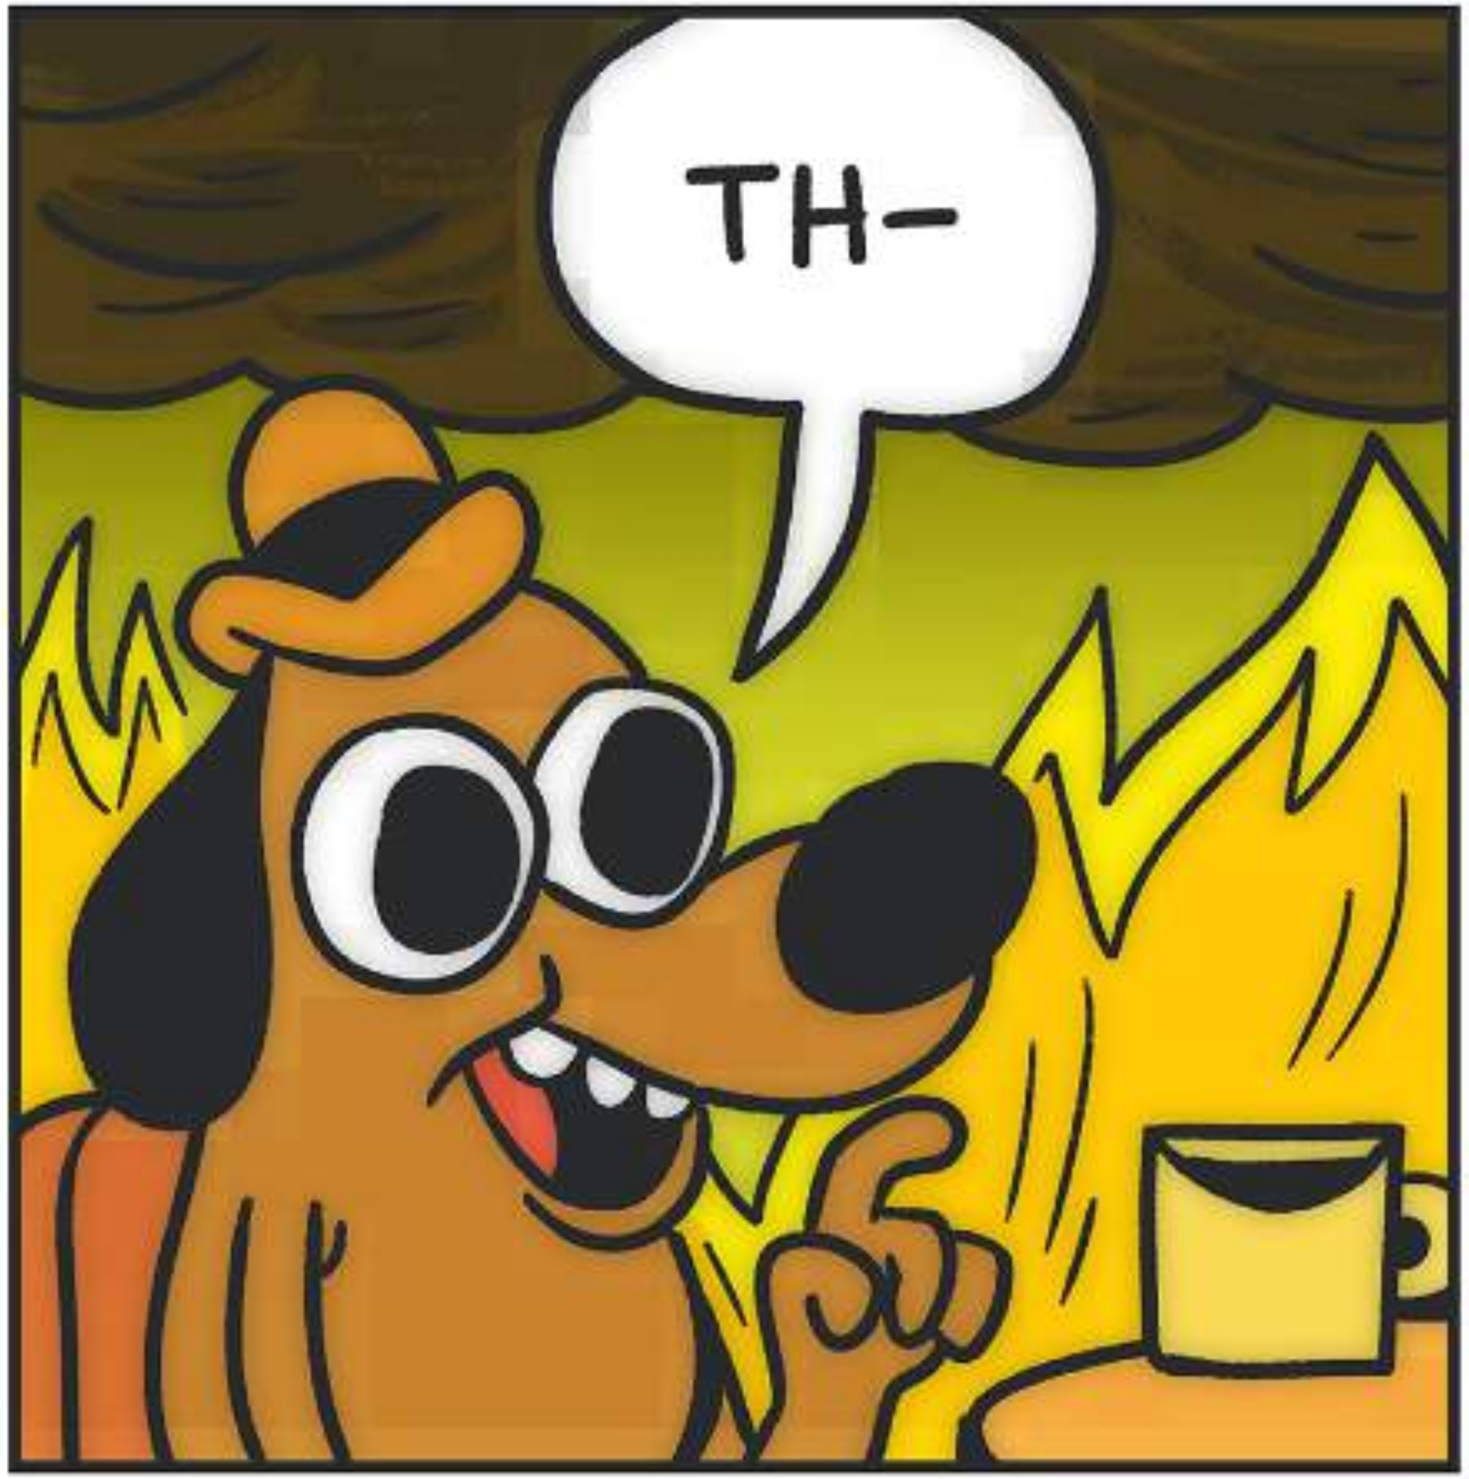
TH-

THIS IS NOT FINE!!

DIABOLICAL
LAUGHTER

Observability

A system is observable when you can ask arbitrary questions about it and receive meaningful answers without having to resort to writing new code or command line tools.

It lets you discover unknown-unknowns and debug in production.

Three Pillars of Observability

- Metrics
- Logs
- APM

These aren't pillars.

Three Pillars of Carpentry?

stahp.

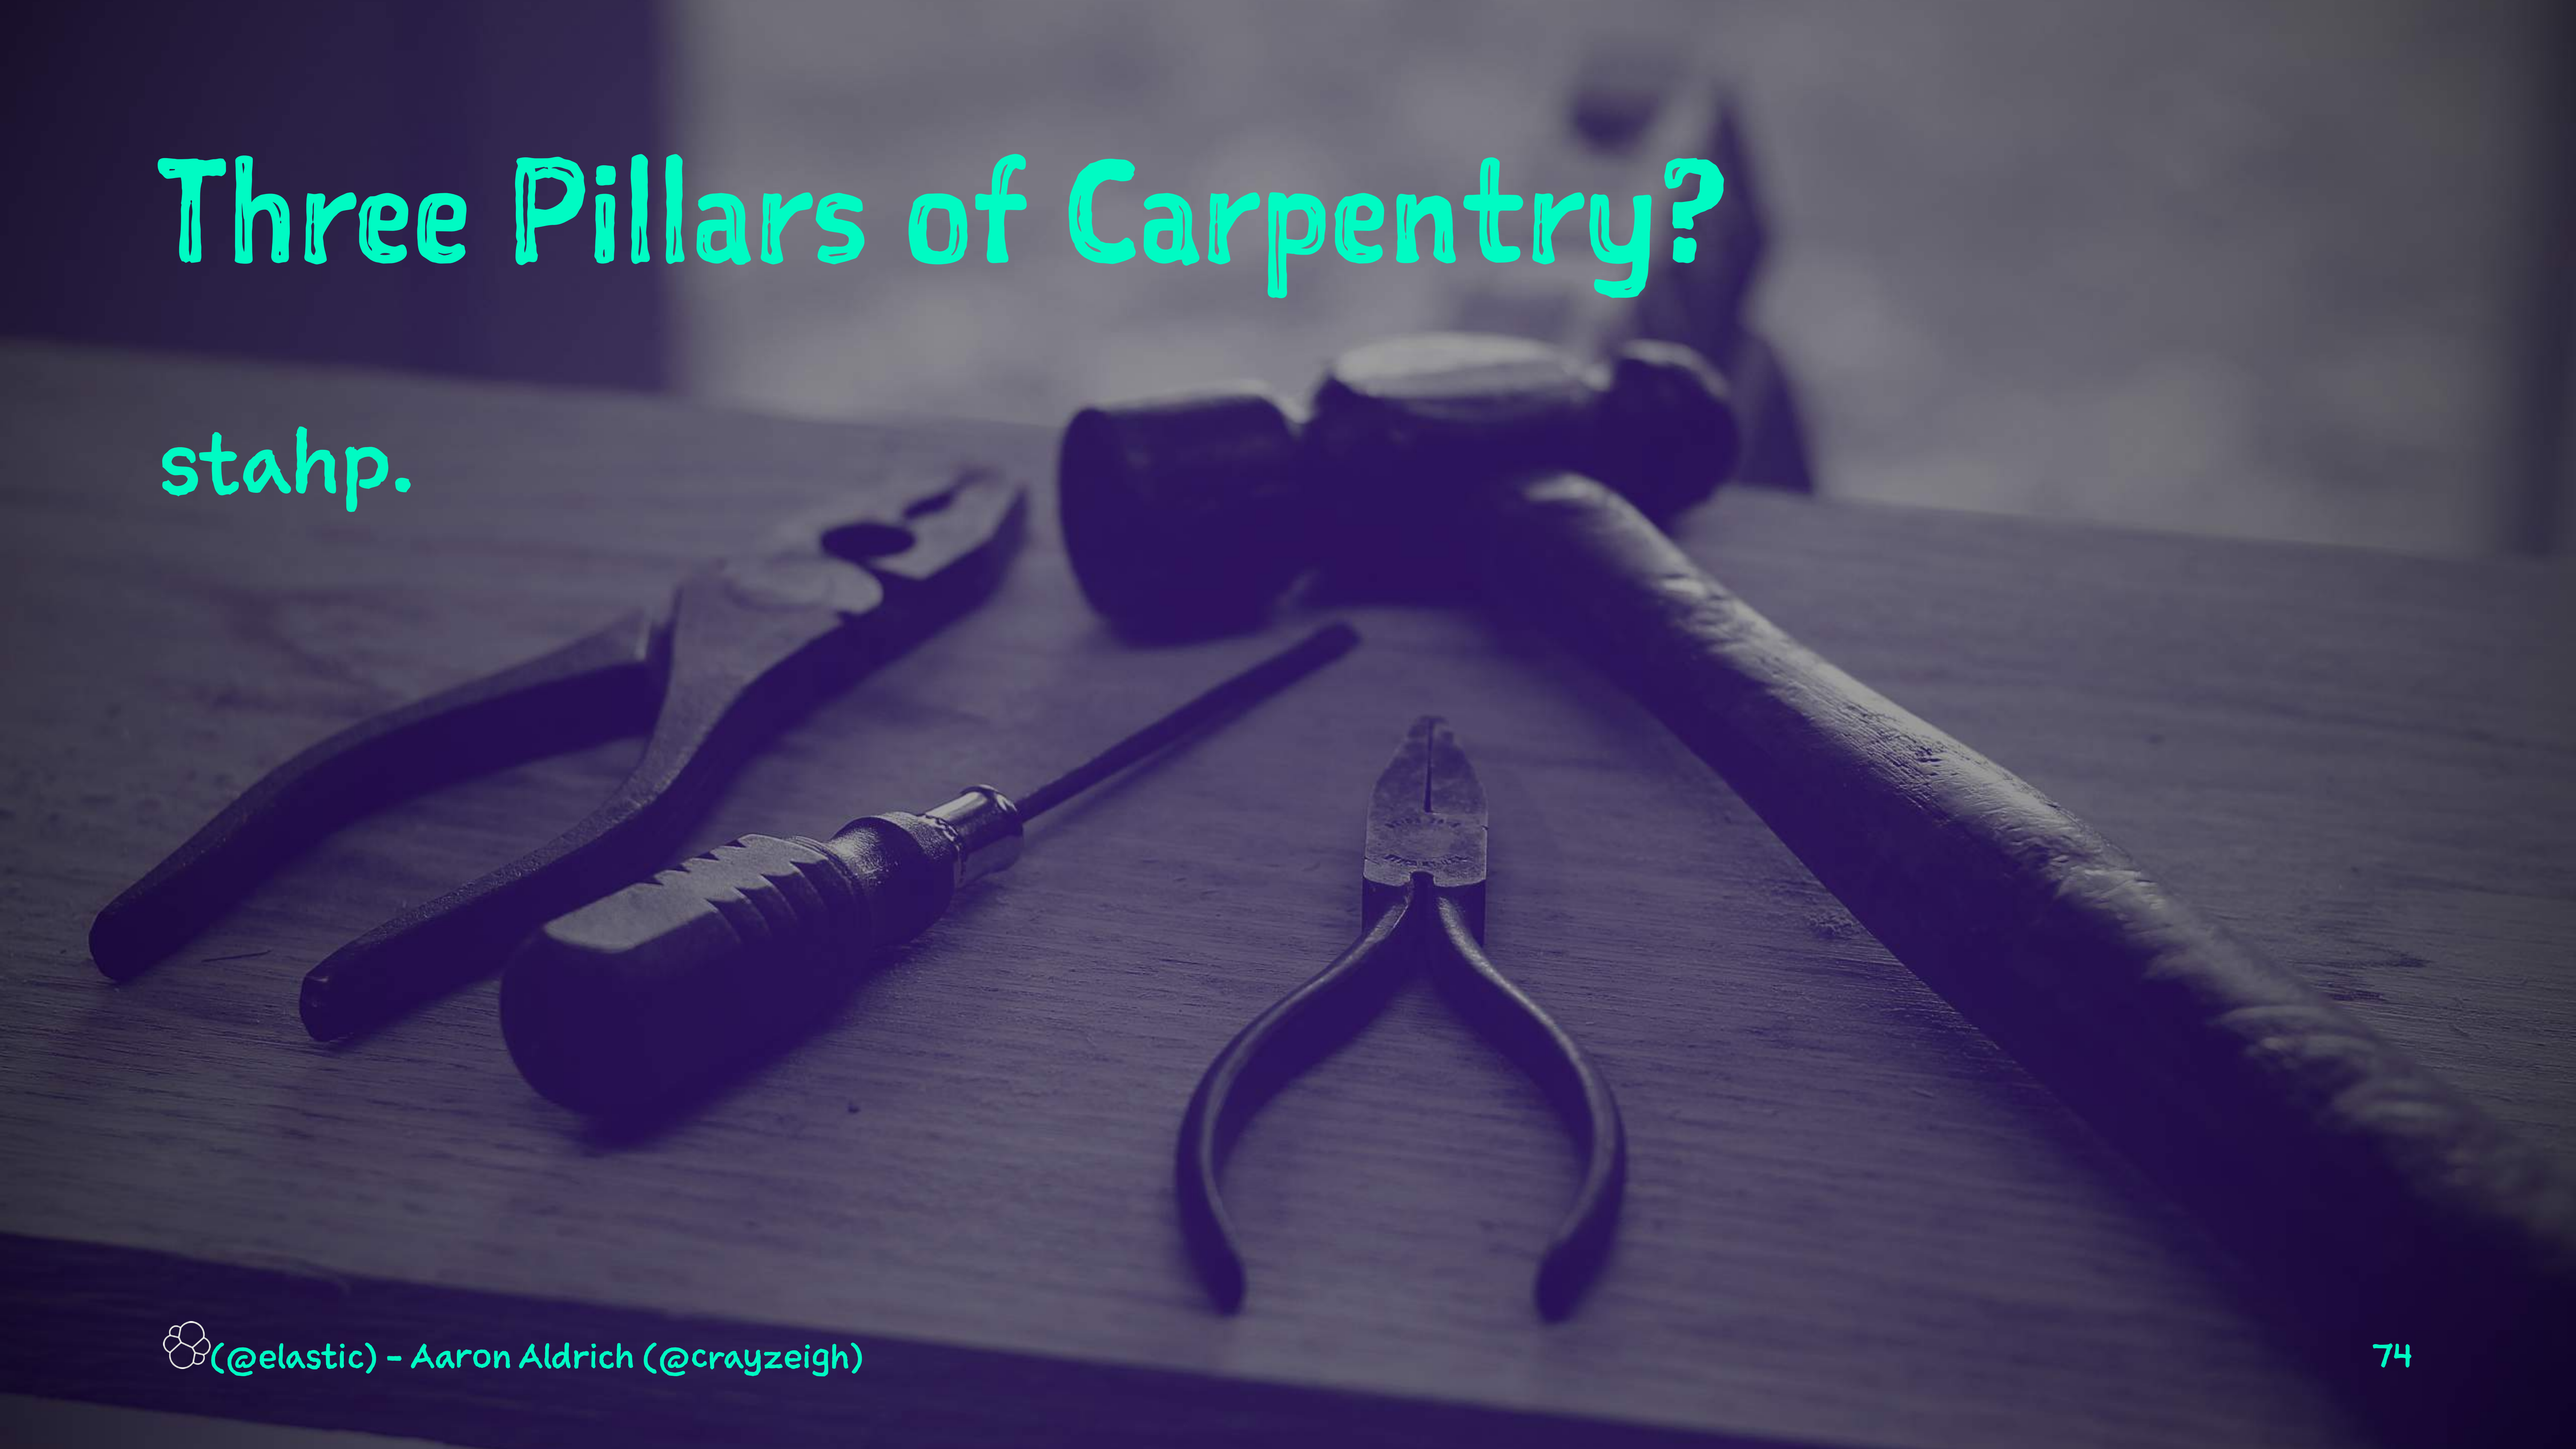

They're tools, not pillars

You need to know how to use them

Metrics

- Great, not on their own
- largely contextless
- need further notation to be valuable (tags)
- Easy to store lots of them
- collection can be a pain📈

📈 Check out Open Metrics! openmetrics.io

High Cardinality Data

- UUIDs
- raw queries
- comments
- firstname, lastname
- PID/PPID
- app ID
- device ID
- build ID
- IP:port
- shopping cart ID
- userid

What's better at carrying Cardinality?

Events!

(Logs)

But please not these:

```
64.242.88.10 - - [07/Mar/2004:16:05:49 -0800] "GET /twiki/bin/edit/Main/Double_bounce_sender?topicparent=Main.ConfigurationVariables HTTP/1.1" 401 12846
64.242.88.10 - - [07/Mar/2004:16:06:51 -0800] "GET /twiki/bin/rdiff/TWiki/NewUserTemplate?rev1=1.3&rev2=1.2 HTTP/1.1" 200 4523
64.242.88.10 - - [07/Mar/2004:16:10:02 -0800] "GET /mailman/listinfo/hsdivision HTTP/1.1" 200 6291
64.242.88.10 - - [07/Mar/2004:16:11:58 -0800] "GET /twiki/bin/view/TWiki/WikiSyntax HTTP/1.1" 200 7352
64.242.88.10 - - [07/Mar/2004:16:20:55 -0800] "GET /twiki/bin/view/Main/DCCAndPostFix HTTP/1.1" 200 5253
64.242.88.10 - - [07/Mar/2004:16:23:12 -0800] "GET /twiki/bin/oops/TWiki/AppendixFileSystem?template=oopsmore¶m1=1.12¶m2=1.12 HTTP/1.1" 200 11382
64.242.88.10 - - [07/Mar/2004:16:24:16 -0800] "GET /twiki/bin/view/Main/PeterThoeny HTTP/1.1" 200 4924
64.242.88.10 - - [07/Mar/2004:16:29:16 -0800] "GET /twiki/bin/edit/Main/Header_checks?topicparent=Main.ConfigurationVariables HTTP/1.1" 401 12851
64.242.88.10 - - [07/Mar/2004:16:30:29 -0800] "GET /twiki/bin/attach/Main/OfficeLocations HTTP/1.1" 401 12851
64.242.88.10 - - [07/Mar/2004:16:31:48 -0800] "GET /twiki/bin/view/TWiki/WebTopicEditTemplate HTTP/1.1" 200 3732
64.242.88.10 - - [07/Mar/2004:16:32:50 -0800] "GET /twiki/bin/view/Main/WebChanges HTTP/1.1" 200 40520
64.242.88.10 - - [07/Mar/2004:16:33:53 -0800] "GET /twiki/bin/edit/Main/Smtpd_etrn_restrictions?topicparent=Main.ConfigurationVariables HTTP/1.1" 401 12851
64.242.88.10 - - [07/Mar/2004:16:35:19 -0800] "GET /mailman/listinfo/business HTTP/1.1" 200 6379
64.242.88.10 - - [07/Mar/2004:16:36:22 -0800] "GET /twiki/bin/rdiff/Main/WebIndex?rev1=1.2&rev2=1.1 HTTP/1.1" 200 46373
64.242.88.10 - - [07/Mar/2004:16:37:27 -0800] "GET /twiki/bin/view/TWiki/DontNotify HTTP/1.1" 200 4140
64.242.88.10 - - [07/Mar/2004:16:39:24 -0800] "GET /twiki/bin/view/Main/TokyoOffice HTTP/1.1" 200 3853
64.242.88.10 - - [07/Mar/2004:16:43:54 -0800] "GET /twiki/bin/view/Main/MikeMannix HTTP/1.1" 200 3686
64.242.88.10 - - [07/Mar/2004:16:45:56 -0800] "GET /twiki/bin/attach/Main/PostfixCommands HTTP/1.1" 401 12846
64.242.88.10 - - [07/Mar/2004:16:47:12 -0800] "GET /robots.txt HTTP/1.1" 200 68
64.242.88.10 - - [07/Mar/2004:16:47:46 -0800] "GET /twiki/bin/rdiff/Know/ReadmeFirst?rev1=1.5&rev2=1.4 HTTP/1.1" 200 5724
64.242.88.10 - - [07/Mar/2004:16:49:04 -0800] "GET /twiki/bin/view/Main/TWikiGroups?rev=1.2 HTTP/1.1" 200 5162
64.242.88.10 - - [07/Mar/2004:16:50:54 -0800] "GET /twiki/bin/rdiff/Main/ConfigurationVariables HTTP/1.1" 200 59679
64.242.88.10 - - [07/Mar/2004:16:52:35 -0800] "GET /twiki/bin/edit/Main/Flush_service_name?topicparent=Main.ConfigurationVariables HTTP/1.1" 401 12851
64.242.88.10 - - [07/Mar/2004:16:53:46 -0800] "GET /twiki/bin/rdiff/TWiki/TWikiRegistration HTTP/1.1" 200 34395
64.242.88.10 - - [07/Mar/2004:16:54:55 -0800] "GET /twiki/bin/rdiff/Main/NicholasLee HTTP/1.1" 200 7235
64.242.88.10 - - [07/Mar/2004:16:56:39 -0800] "GET /twiki/bin/view/Sandbox/WebHome?rev=1.6 HTTP/1.1" 200 8545
64.242.88.10 - - [07/Mar/2004:16:58:54 -0800] "GET /mailman/listinfo/administration HTTP/1.1" 200 6459
lordgun.org - - [07/Mar/2004:17:01:53 -0800] "GET /razor.html HTTP/1.1" 200 2869
64.242.88.10 - - [07/Mar/2004:17:09:01 -0800] "GET /twiki/bin/search/Main/SearchResult?scope=text®ex=on&search=Joris%20*Benschop[^A-Za-z] HTTP/1.1" 200 4284
```

Structured Data🕶

```
{
  "message":"user_deleted",
  "user": {
    "id":6,
    "email":"crayzeigh@example.com",
    "created_at":"2015-12-11T04:31:46.828Z",
    "updated_at":"2015-12-11T04:32:18.340Z",
    "name":"crayzeigh",
    "role":"user",
    "invitation_token":null,
    "invitation_created_at":null,
    "invitation_sent_at":null,
    "invitation_accepted_at":null,
    "invitation_limit":null,
    "invited_by_id":null,
    "invited_by_type":null,
    "invitations_count":0
  },
  "@timestamp":"2015-12-11T13:35:50.070+00:00",
  "@version":"1",
  "severity":"INFO",
  "host":"app1-web1",
  "type":"apps"
}
```

🕶 from James Turnbull: https://www.kartar.net/2015/12/structured-logging/

Generate LOTS of events use sampling to store them

OK let's talk about APM

APM / opbeans-node / Transactions / GET /api/stats

GET /api/stats

Search transactions and errors... (E.g. transaction.duration.us > 300000 AND context.response.status_code >= 400)

Last 18 hours

Distributed Tracing 🔍

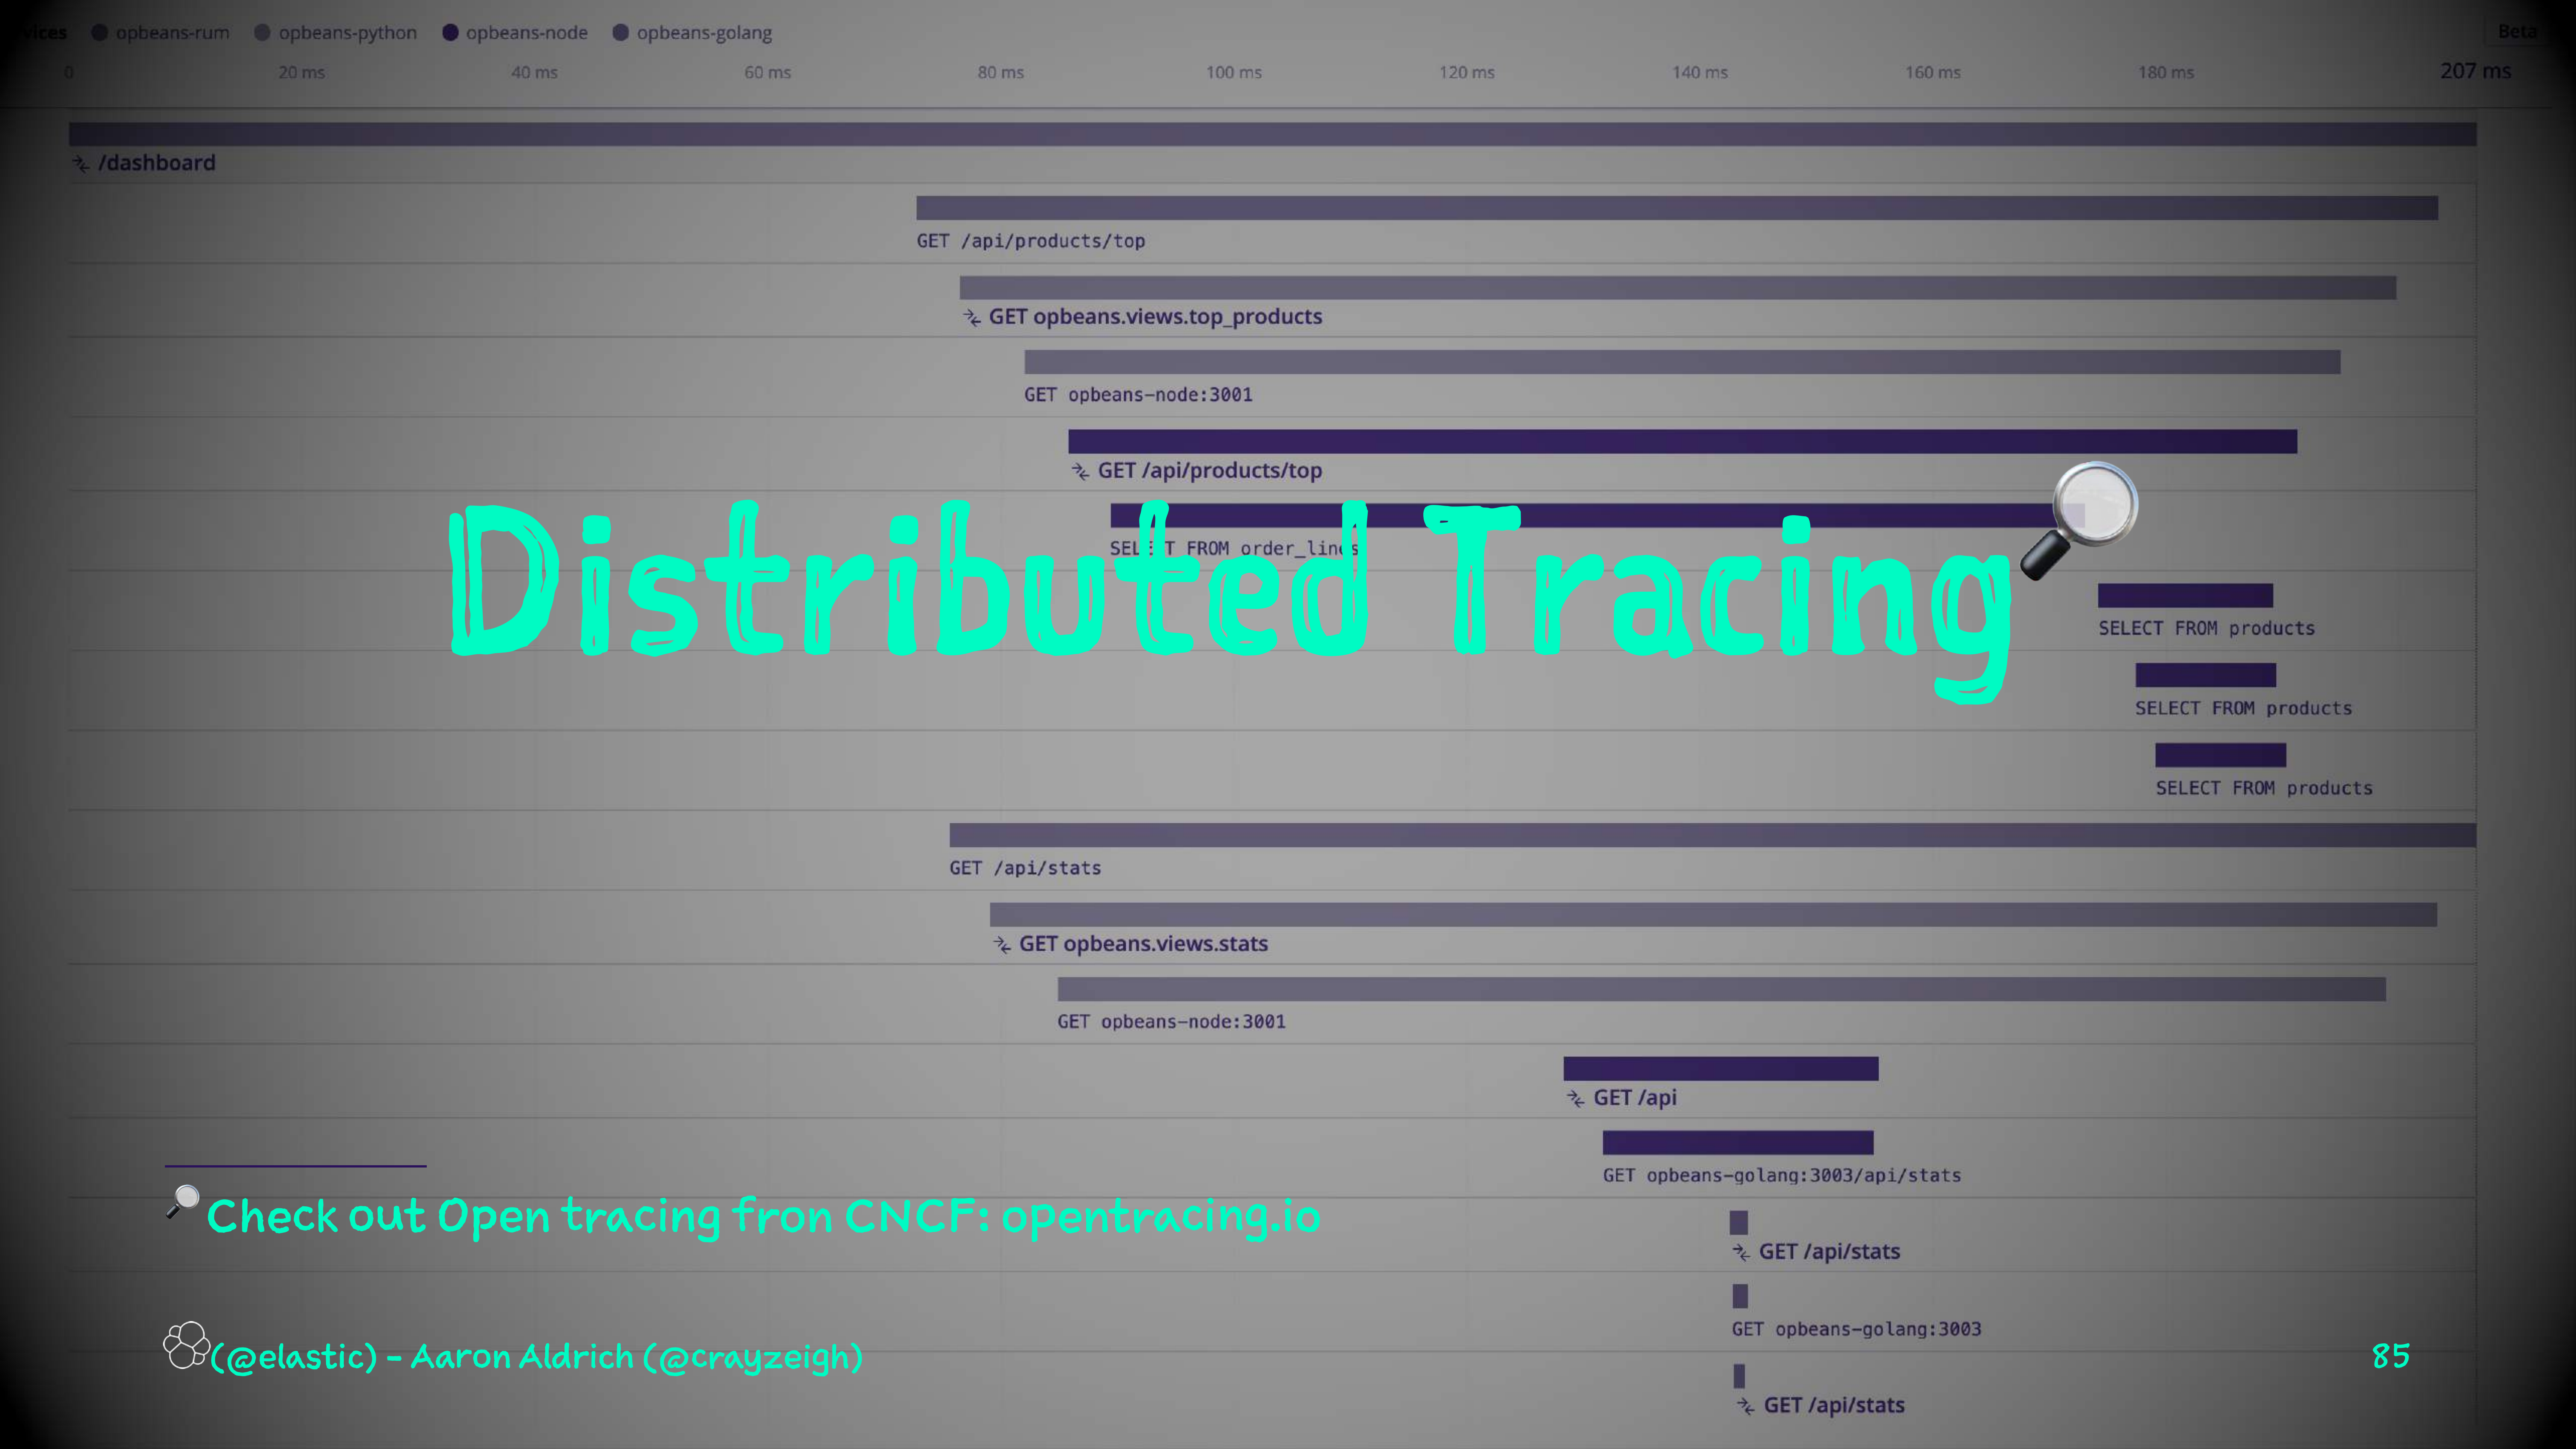

```
vices ● opbeans-rum ● opbeans-python ● opbeans-node ● opbeans-golang                Beta
0    20 ms    40 ms    60 ms    80 ms    100 ms    120 ms    140 ms    160 ms    180 ms    207 ms

/dashboard
        GET /api/products/top
         GET opbeans.views.top_products
          GET opbeans-node:3001
           GET /api/products/top
            SELECT FROM order_lines
                                      SELECT FROM products
                                       SELECT FROM products
                                        SELECT FROM products
        GET /api/stats
         GET opbeans.views.stats
          GET opbeans-node:3001
               GET /api
               GET opbeans-golang:3003/api/stats
                 GET /api/stats
                 GET opbeans-golang:3003
                 GET /api/stats
```

🔎 Check out Open tracing fron CNCF: opentracing.io

Instrumentation:

SLIs are a good place to start

Kill Staging:

Test in Production

This doesn't eliminate QA or testing

(please test before prod)

Kill your staging environment

- always out of sync
- can't replicate prod traffic anyway
- definitely can't replicate real users
- replace with feature flags and canary deploys🚀

🚀 Launch Darkly talks about this a lot. You should listen to what they have to say.

O11y ❤️'s QA

Start leveraging a common toolset

Every Dashboard sucks

Not really, some dashboards are pretty good

It's about Storytelling

know your audience

Ops & Incident Response

- Interactive
- Iterative
- Involve search bars

Vendor Warning:

Search &

Common Data Schema

Making O11y Evangelists

Don't just start making changes

History is important

Change conducted poorly breaks organizations

top-down mandated change never works☠️

☠️Did you know "defenestration" is the act of throwing someone out a window?

Talk to other parts of the business to understand what stories they value

LISTEN

It's all about context

Start measuring business values

Who else might care about dashboards?

What data can we expose to the rest of the business?

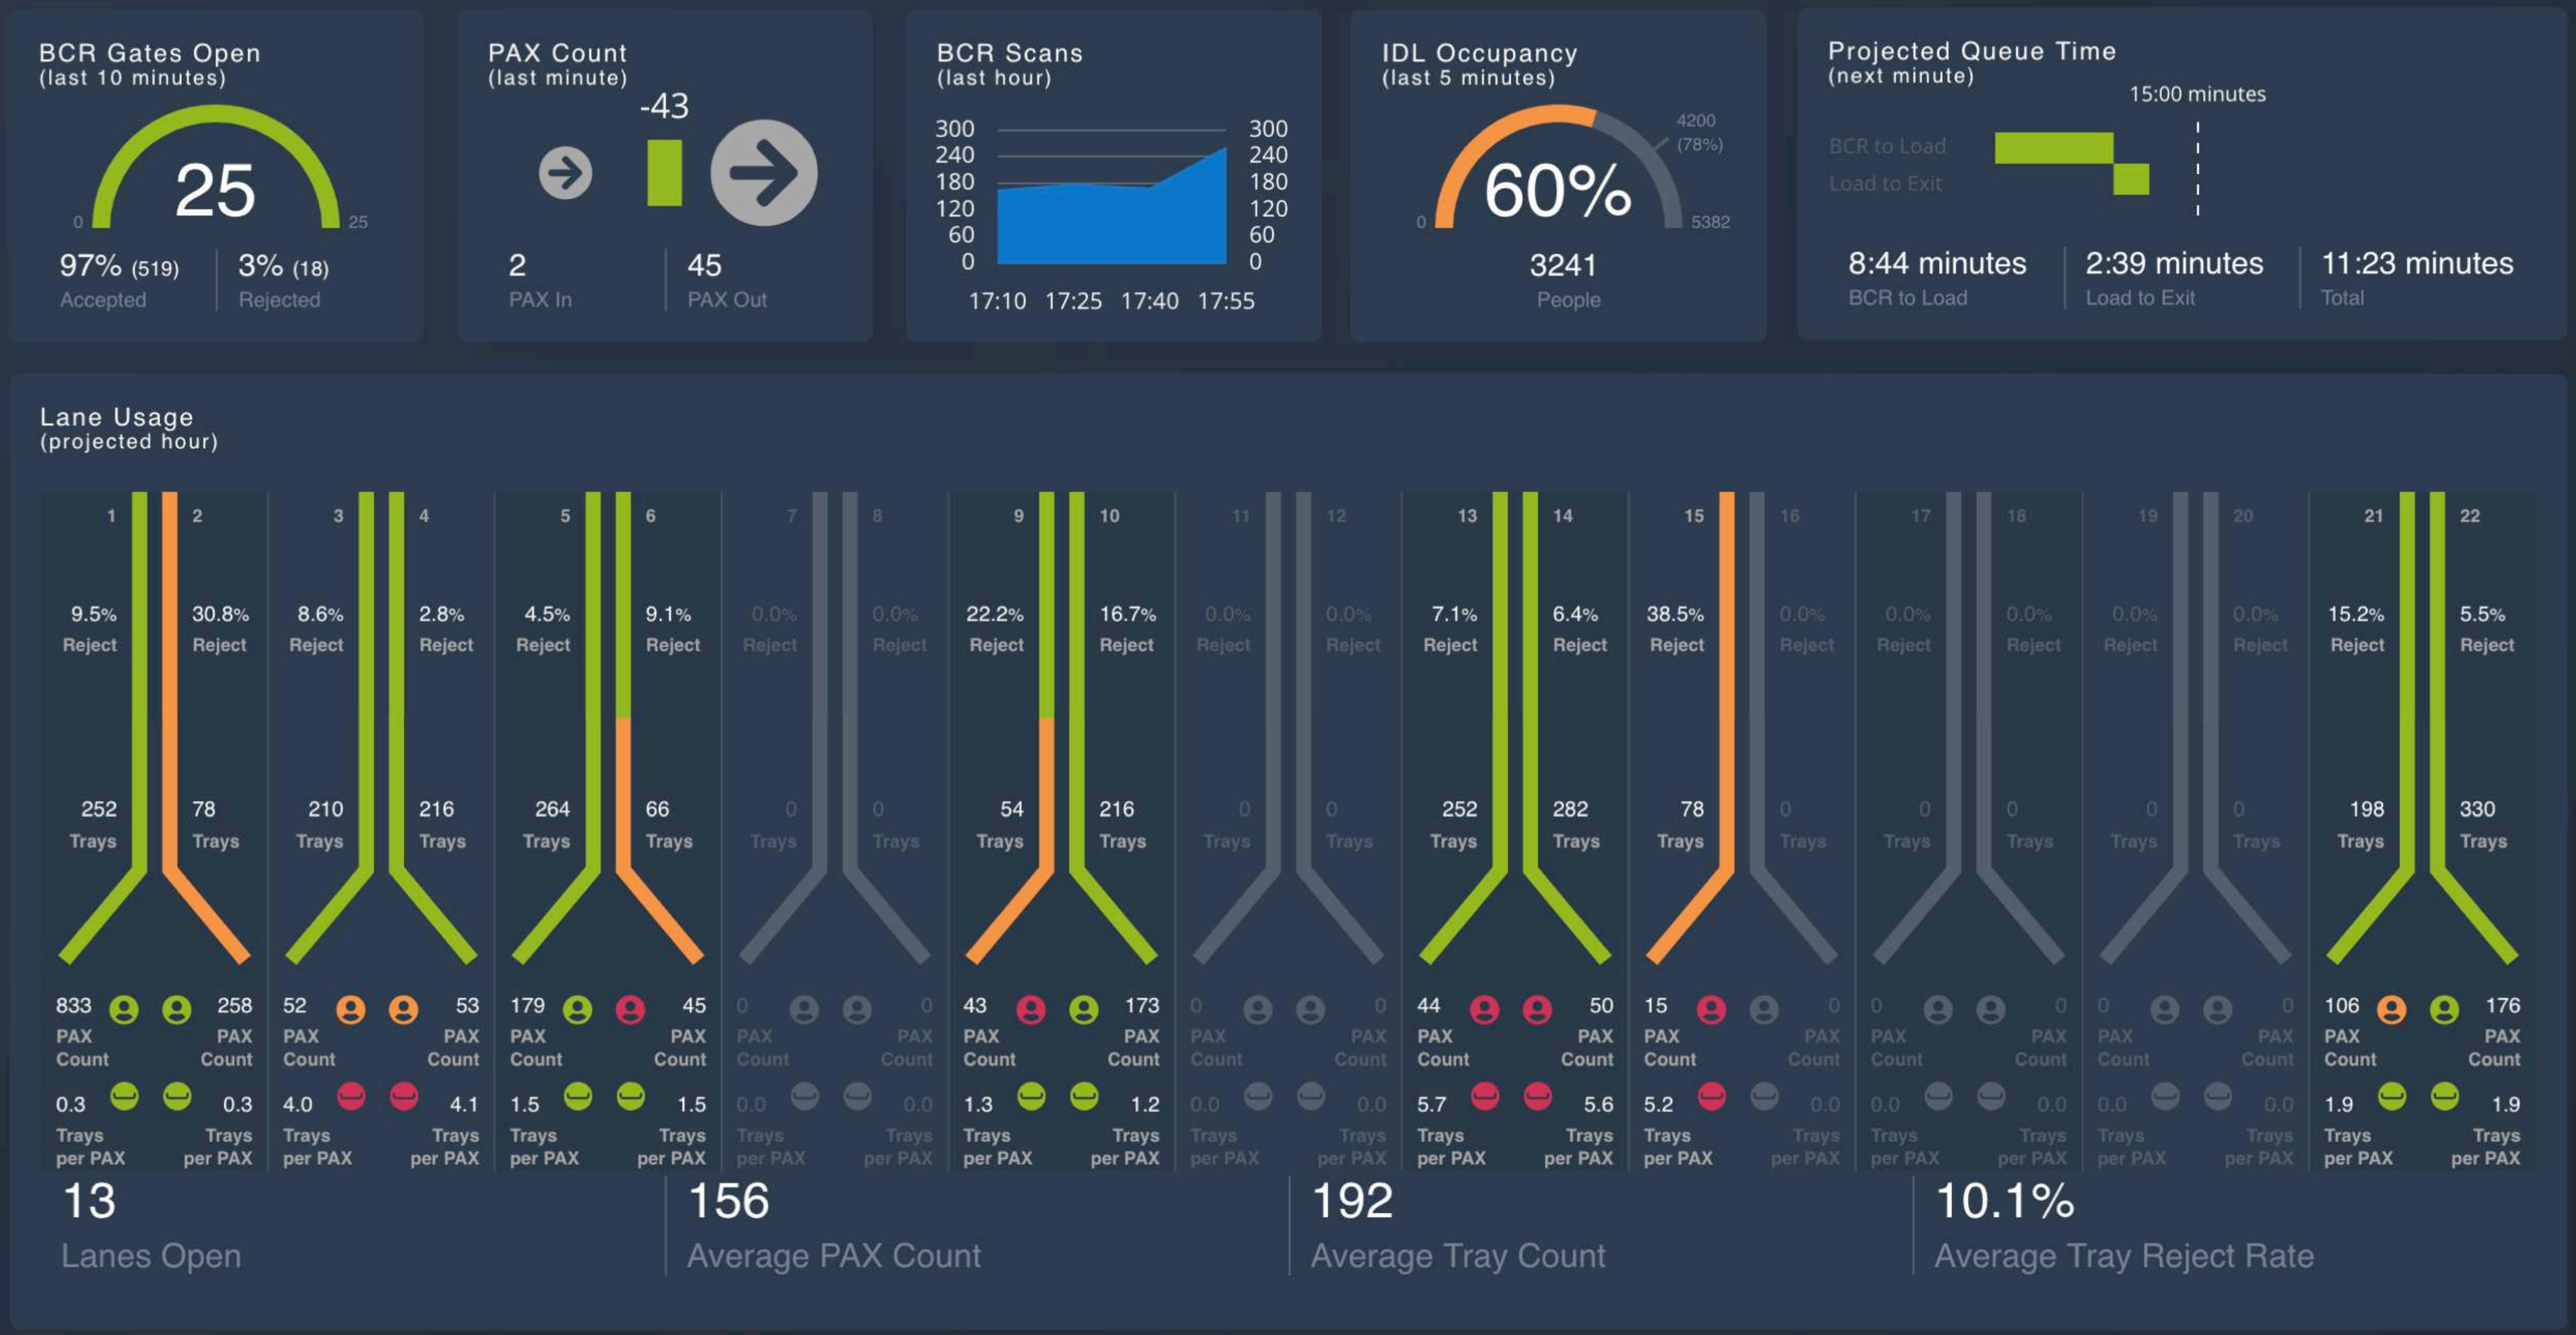

BCR Gates Open (last 10 minutes)

25

97% (519) Accepted | 3% (18) Rejected

PAX Count (last minute)

2 PAX In | 45 PAX Out

BCR Scans (last hour)

IDL Occupancy (last 5 minutes)

60%

3241 People

Projected Queue Time (next minute)

BCR to Load

Load to Exit

8:44 minutes BCR to Load | 2:39 minutes Load to Exit | 11:23 minutes Total

Lane Usage (projected hour)

Lane	Reject	Trays	PAX Count	Trays per PAX
1	9.5%	252	833	0.3
2	30.8%	78	258	0.3
3	8.6%	210	52	4.0
4	2.8%	216	53	4.1
5	4.5%	264	179	1.5
6	9.1%	66	45	1.5
7	0.0%	0	0	0.0
8	0.0%	0	0	0.0
9	22.2%	54	43	1.3
10	16.7%	216	173	1.2
11	0.0%	0	0	0.0
12	0.0%	0	0	0.0
13	7.1%	252	44	5.7
14	6.4%	282	50	5.6
15	38.5%	78	15	5.2
16	0.0%	0	0	0.0
17	0.0%	0	0	0.0
18	0.0%	0	0	0.0
19	0.0%	0	0	0.0
20	0.0%	0	0	0.0
21	15.2%	198	106	1.9
22	5.5%	330	176	1.9

13 Lanes Open | 156 Average PAX Count | 192 Average Tray Count | 10.1% Average Tray Reject Rate

The information in this dashboard is sample data only

MAG London Stansted Airport

Digital Marketing

Marketing Funnel

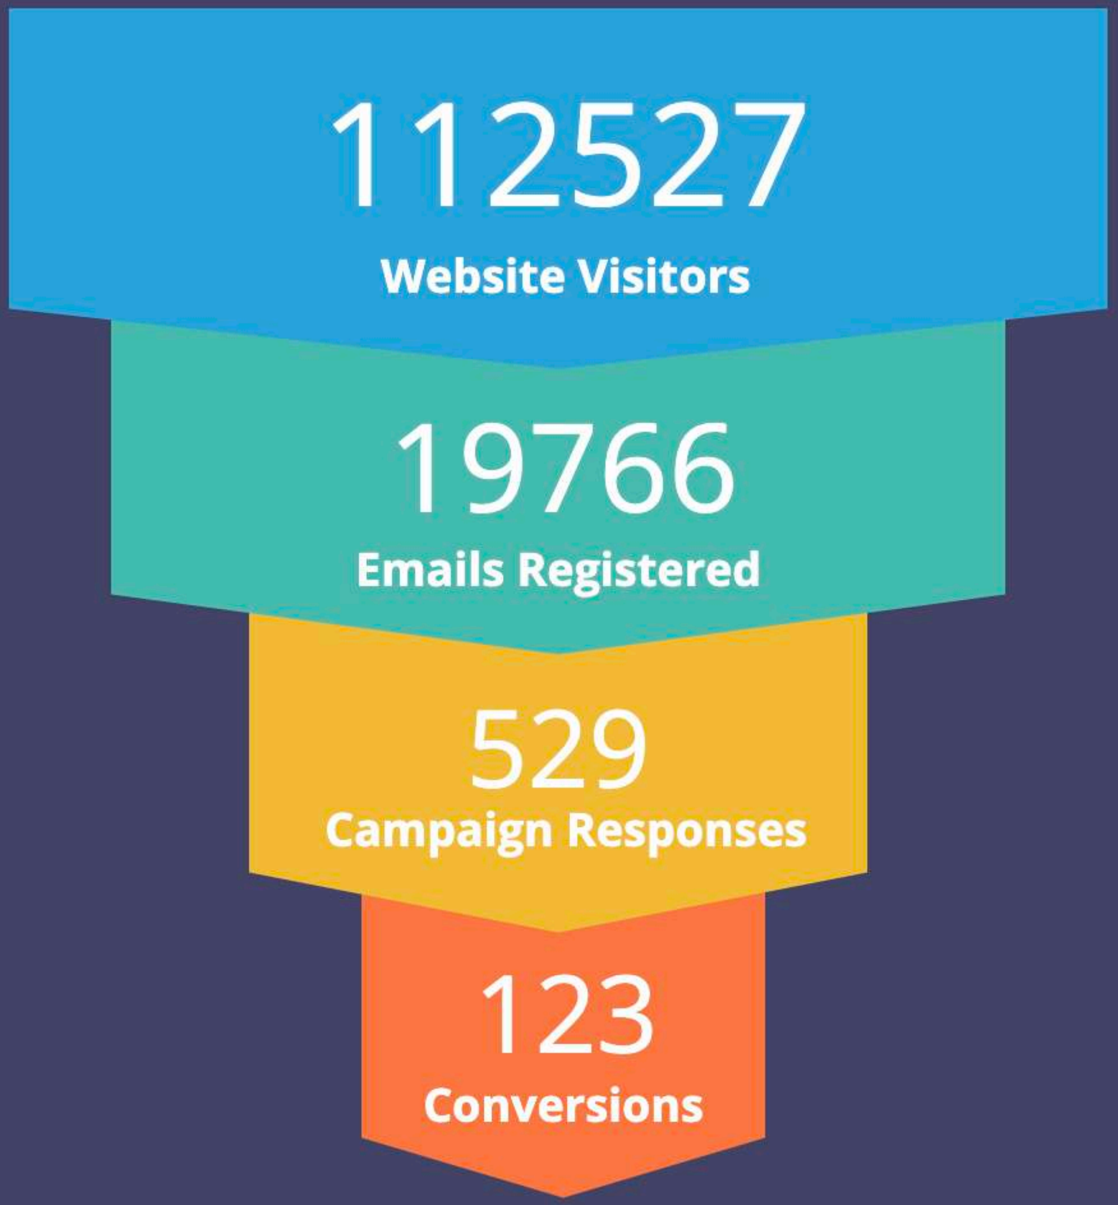
112527
Website Visitors
19766
Emails Registered
529
Campaign Responses
123
Conversions

Conversion Ratio

1.03%

Bounce Rate

49.4%

Subscribers by Platform

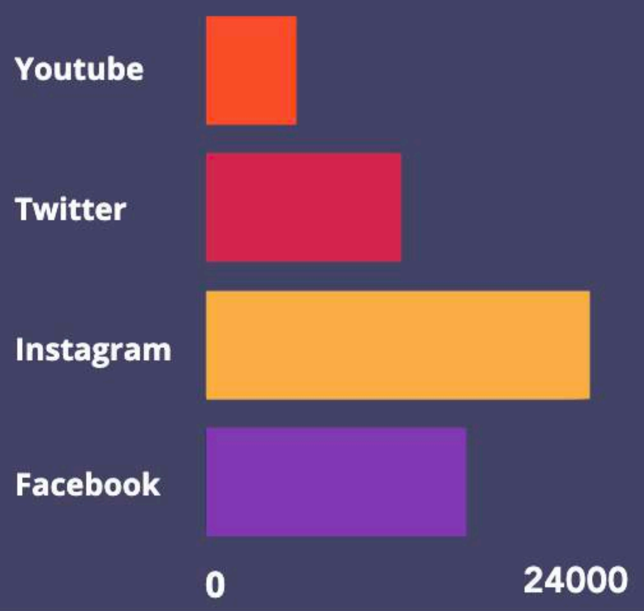

Email Campaigns

24

Newsletter Open Rate

53.8%

Campaign Responses

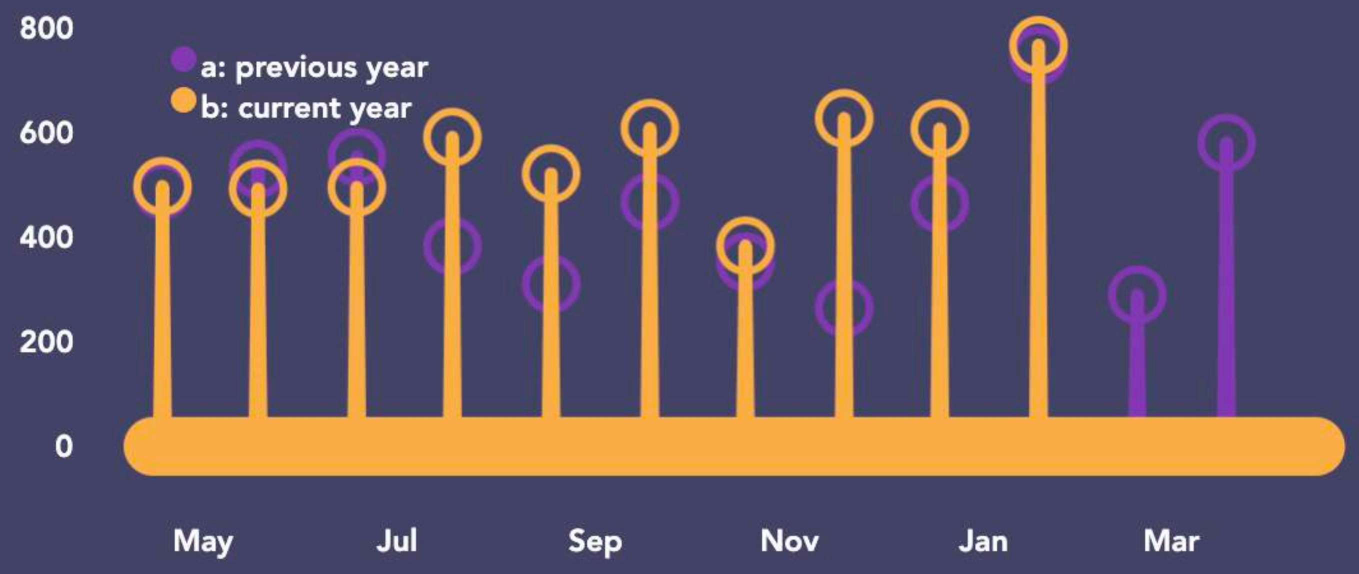

Weekly Visitors by Day

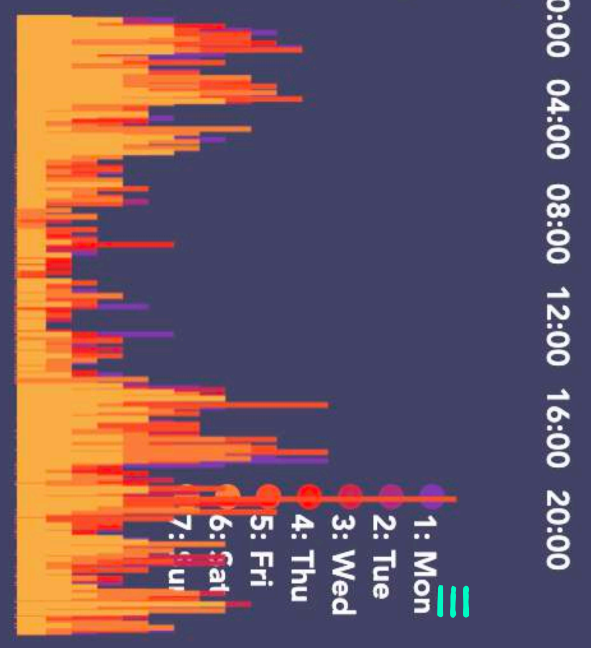

$351k

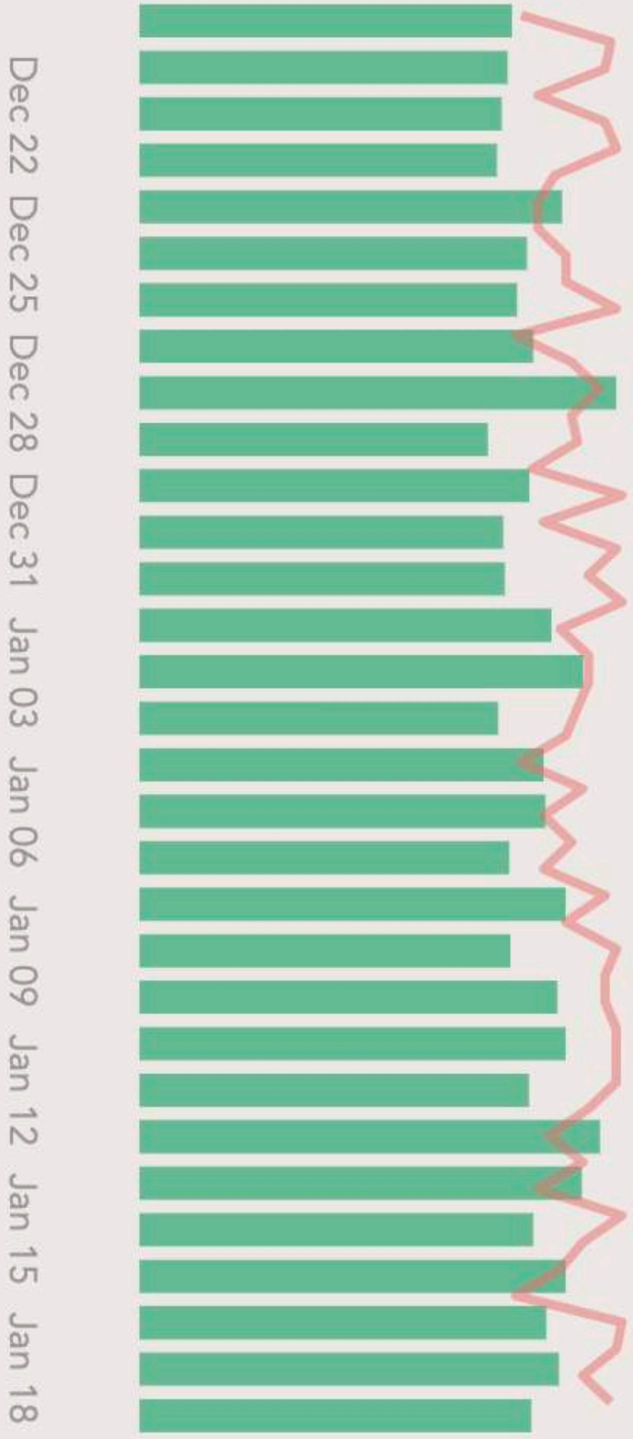
Dec 22
Dec 25
Dec 28
Dec 31
Jan 03
Jan 06
Jan 09
Jan 12
Jan 15
Jan 18

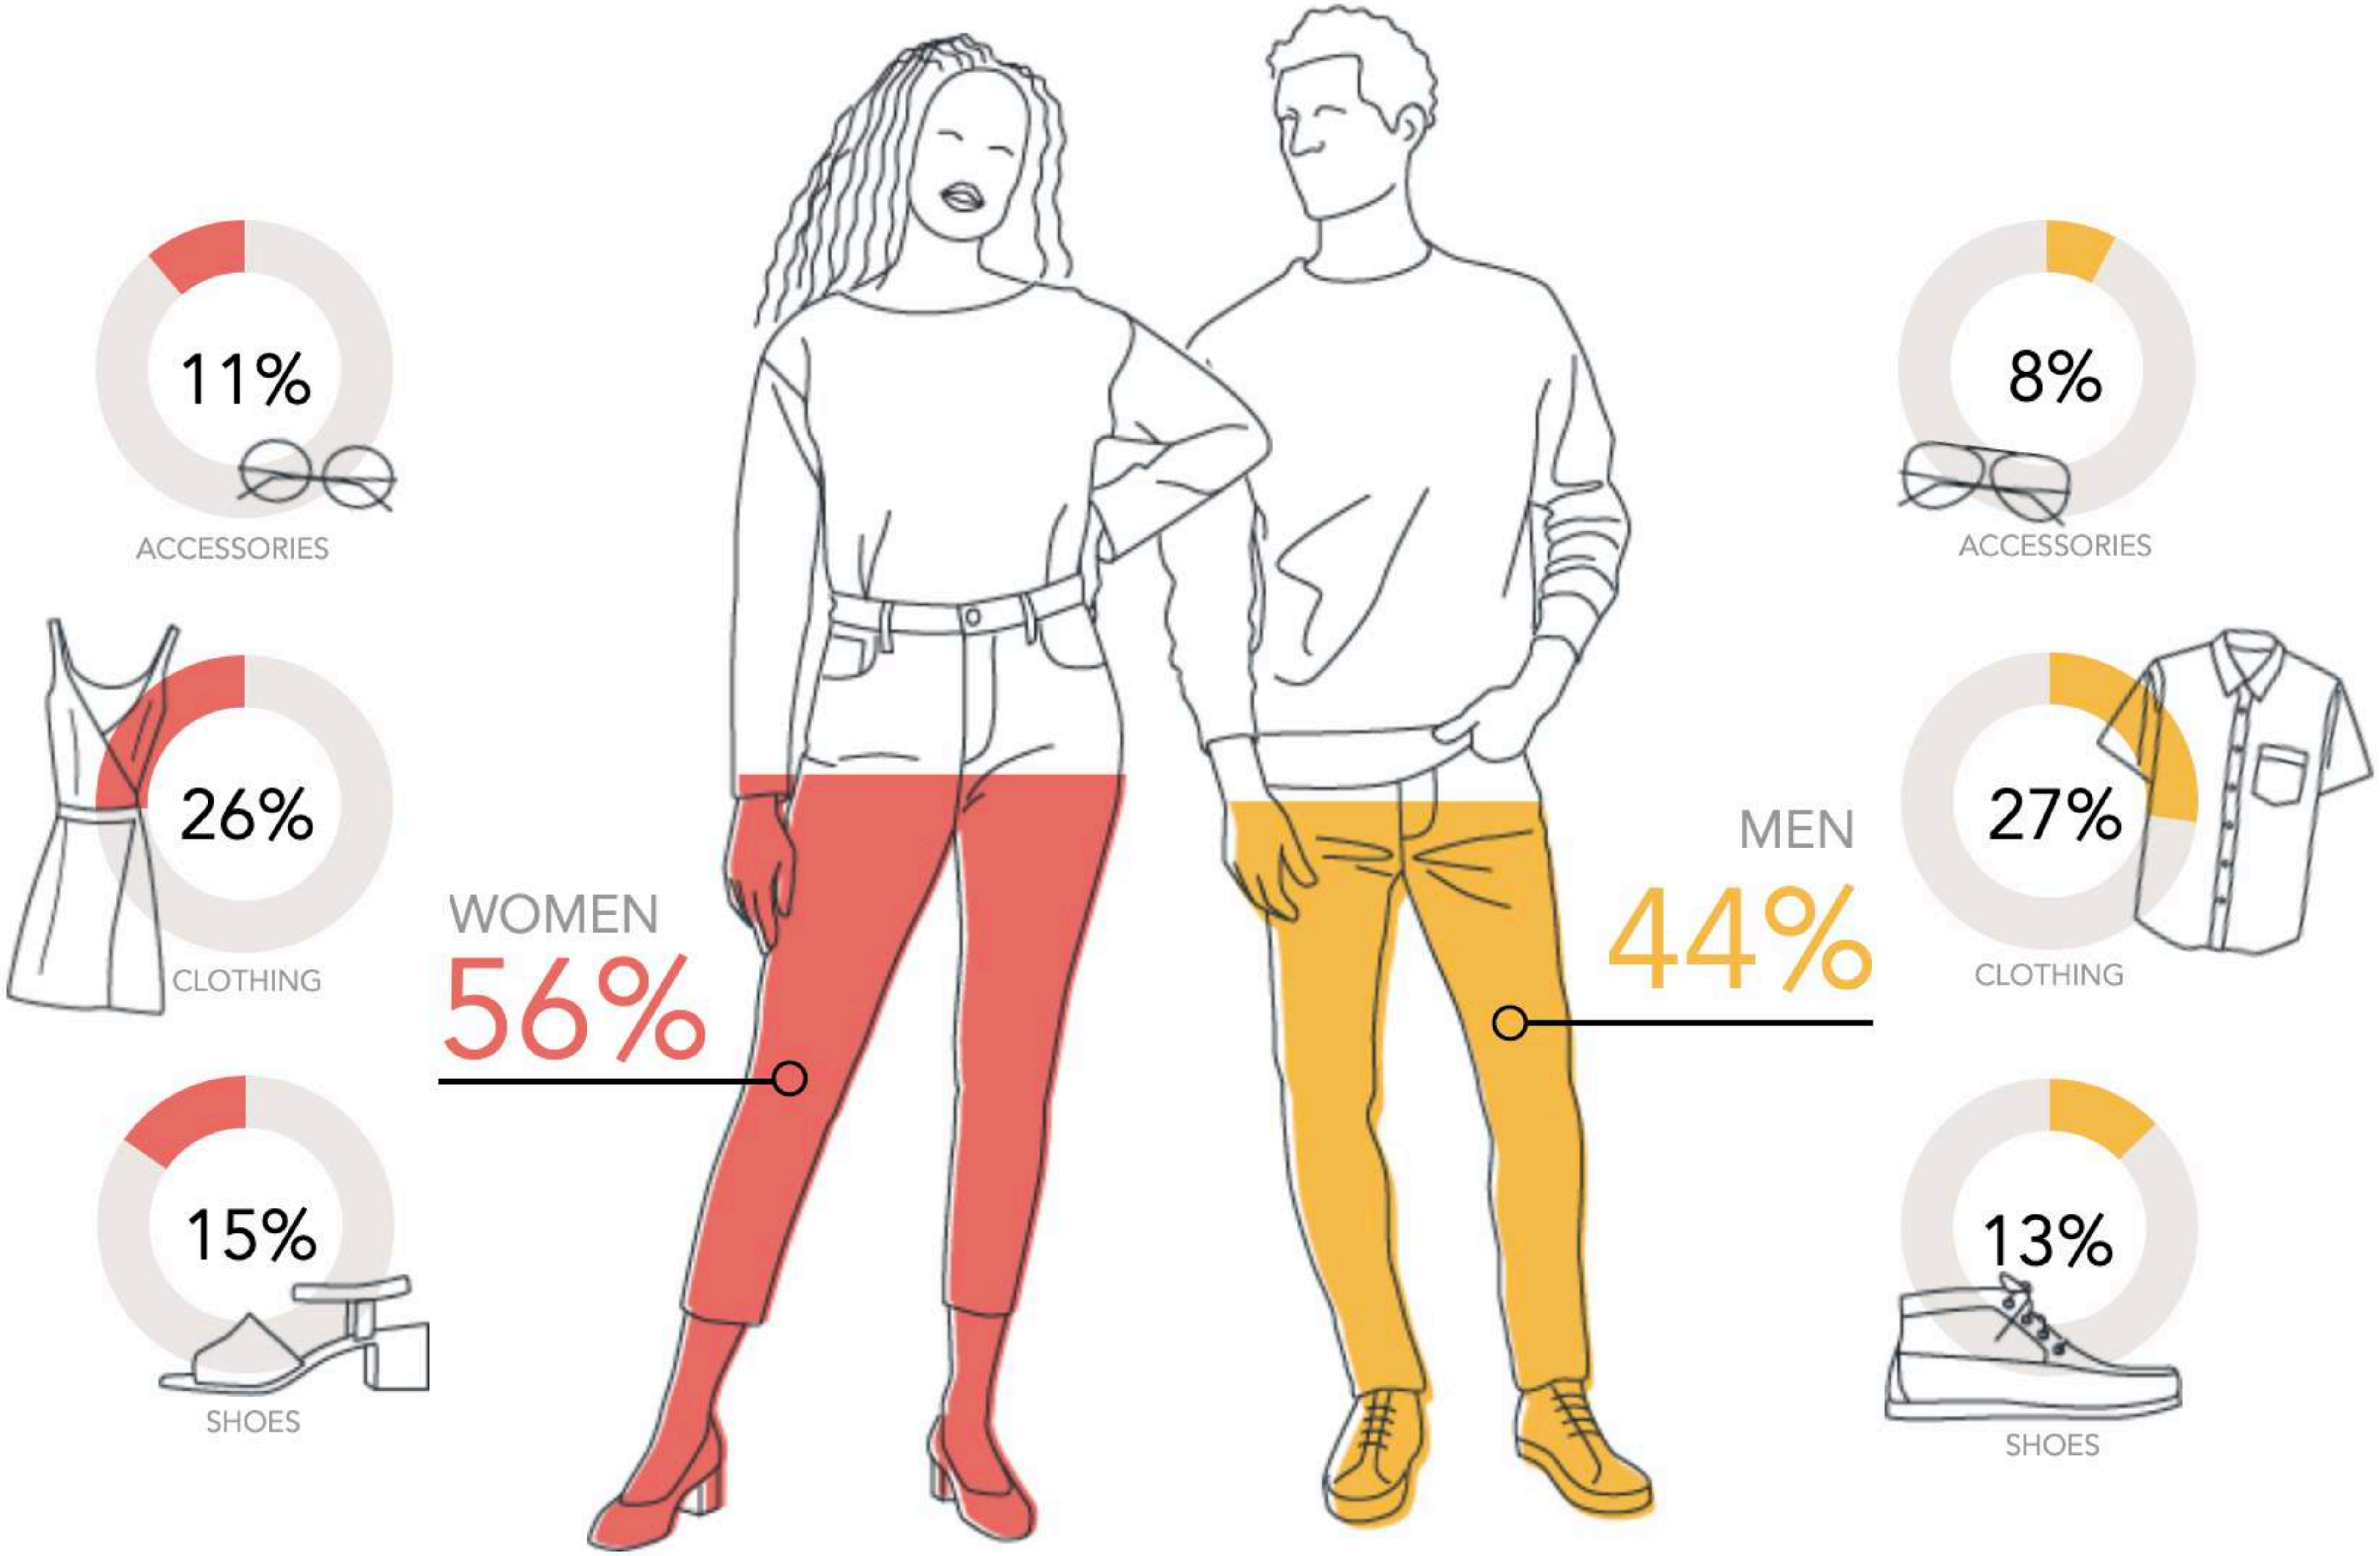
11%
ACCESSORIES
26%
CLOTHING
15%
SHOES
WOMEN
56%
MEN
44%
8%
ACCESSORIES
27%
CLOTHING
13%
SHOES

Dashboards help tell stories with context

Share results

Good and Bad

Are your systems up?

Are they responding acceptably?

Who cares?

Are your services delivering value?