

Restoring Confidence in Microservices: Tracing That's More Than Traces

Ben Sigelman (LightStep co-founder, OpenTracing co-creator)

QCon London, March 6, 2019

A person is seen from behind, climbing a steep, reddish-brown rock face. They are wearing a blue long-sleeved shirt, dark pants, and a red and blue climbing harness. They are using their hands and feet to grip the rock, with one hand reaching up to a small opening. A rope is visible running down the side of the cliff. The rock face is textured with various shades of brown and orange, and some small green plants are visible on the right side. The word "confidence" is written in white lowercase letters at the bottom of the image.

confidence



confidence

```
matt@shadowfax:~/Work/Dev/goconvey/src/github.com/smartystreets/goconvey/examples]
```

```
$ go test -v
```

```
=== RUN TestAssertions
```

All assertions should be accessible

Equality assertions should be accessible ✓✓✓✓✓✓✓✓

Numeric comparison assertions should be accessible ✓✓✓✓✓

Container assertions should be accessible ✓✓✓✓

String assertions should be accessible ✓✓✓✓✓✓✓✓

Panic recovery assertions should be accessible ✓✓✓✓

Type-checking assertions should be accessible ✓✓

Time assertions should be accessible ✓✓✓✓✓✓✓✓

44 assertions thus far

```
--- PASS: TestAssertions (0.04 seconds)
```

```
=== RUN TestScoring
```

Subject: Bowling Game Scoring

Given a fresh score card

When all gutter balls are thrown

The score should be zero ✓

When all throws knock down only one pin

The score should be 20 ✓

When a spare is thrown

The score should include a spare bonus. ✓

When a strike is thrown

The score should include a strike bonus. ✓

When all strikes are thrown

The score should be 300. ✓

confidence in software

Cluster Health

Cluster Pod Usage



Cluster CPU Usage




Cluster Memory Usage



Cluster Disk Usage



Cluster Pod Capacity




Cluster CPU Capacity



Cluster Mem Capacity



Cluster Disk Capacity



Deployments

Deployment Replicas - Up To Date

Time	Metric ▾	Value
2019-01-09 23:24:57	up-to-date	1
2019-01-09 23:24:57	desired	1
2019-01-09 23:24:57	ready	1
2019-01-09 23:24:57	available	1

Deployment Replicas

24

Deployment Replicas - Updated

24

Deployment Replicas - Unavailable

1

Node

Number Of Nodes

10

Nodes Out of Disk

Nodes Unavailable

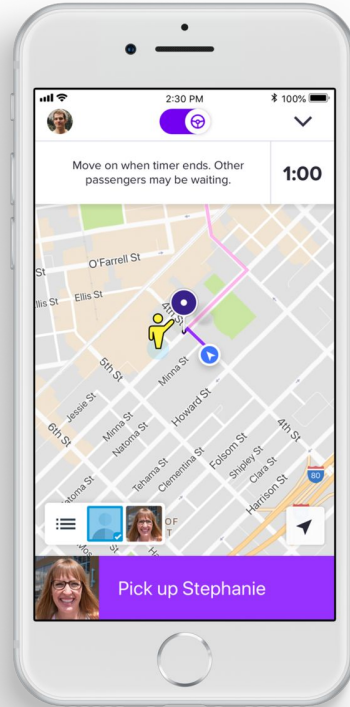
confidence in software

Leverage and Complexity

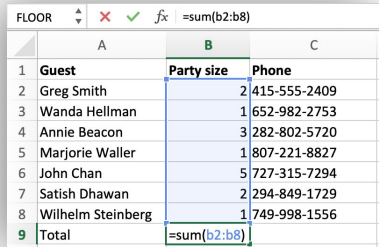
Software with small leverage

FLOOR ⬆ ✖ ✓ <i>fx</i> =sum(b2:b8)			
	A	B	C
1	Guest	Party size	Phone
2	Greg Smith	2	415-555-2409
3	Wanda Hellman	1	652-982-2753
4	Annie Beacon	3	282-802-5720
5	Marjorie Waller	1	807-221-8827
6	John Chan	5	727-315-7294
7	Satish Dhawan	2	294-849-1729
8	Wilhelm Steinberg	1	749-998-1556
9	Total	=sum(b2:b8)	


Software with large leverage



Software: a continuum



	A	B	C
1	Guest	Party size	Phone
2	Greg Smith	2	415-555-2409
3	Wanda Hellman	1	652-982-2753
4	Annie Beacon	3	282-802-5720
5	Marjorie Waller	1	807-221-8827
6	John Chan	5	727-315-7294
7	Satish Dhawan	2	294-849-1729
8	Wilhelm Steinberg	1	749-998-1556
9	Total	=sum(b2:b8)	



Small leverage


Large leverage

Clumsy UX

100 UX

Developers: 





Monthly Release Spreadsheets (via Sarah Wells)

QCon London 2019: Mature microservices and how to operate them

27	26	Release	AHP Buildlife 90271 - FT-CMS (replaces 94246)	Pub Support
28	27	Release	AHP Buildlife 90319 - FT-EEDS (replaces 81488)	Pub Support
29	28	Release	AHP Buildlife 90374 - contentAPI (replaces 79050)	Pub Support
30	29	Manual Action	On Prod Apache caches do the following: fweeb20-uviv-uk-p fweeb21-uviv-uk-p fweeb20-uvnj-us-p fweeb21-uvnj-us-p "rm /data/apache/cache/www.ft.com/Q/uz2TBLkvhnyPGL85j6ag.*"	Pub support
31	30	Manual Action	Historical secondaries 1) Connect to ft-cms staging DB machine with ssh. 2) Get the script that is located here: http://anhill.svc.ft.com/viewvc/repoc/online/website/ft-cms/branches/RELEASE-2.17/config/scripts/use_secondaries.sh?revision=120557&view=co and run it with the given command: ./use_secondaries.sh fflaps01-01-uvpr-uk-qa.osb.ft.com 8080 PREDITOR_USERNAME PREDITOR_PASSWORD > use_secondaries.log	Pub support
32	31	Manual Action	Switch the Content API Tomcat services to point directly to FT-CMS Tomcat rather than FT-CMS Apache. This is done via JMX configuration. Use VisualVM, JConsole or equivalent to connect to each host on port 8533. MBeans: com.ft.deliveryplatform->ClientConfiguration->Attributes Change the PageServiceUrl from http://ftcomrender-web-data.svc.ft.com:80/cms/platform/rest/pages to http://flaps20-uviv-uk-p:8080/cms/platform/rest/pages (for the two UK content API tomcats) http://flaps20-uvnj-us-p:8080/cms/platform/rest/pages (for the two US content API tomcats) Change the ContentItemServiceUrl from http://ftcomrender-web-data.svc.ft.com:80/cms/platform/rest/items to http://flaps20-uviv-uk-p:8080/cms/platform/rest/items (for the two UK content API tomcats) http://flaps20-uvnj-us-p:8080/cms/platform/rest/items (for the two US content API tomcats)	Pub Support

Overview Deploy FT-CMS Rollback FT-CMS Content API PVCs Test PVCs Rollback Test PVCs


Sarah Wells March 04, 2019

Technology 3 490

Crisis of Confidence #1


Velocity (or lack thereof)

Org design, human comms, and microservices



You will inevitably ship your org chart

Déjà vu...



Crystal Hirschorn

@cfhirschorn



"You will inevitably ship your org chart" aka Conway's Law.
Love this illustration from .@el_bhs #QConLondon

♡ 7 3:21 PM - Mar 4, 2019




See Crystal Hirschorn's other Tweets



Regaining confidence in velocity

	A	B	C
1	Guest	Party size	Phone
2	Greg Smith	2	415-555-2409
3	Wanda Hellman	1	652-982-2753
4	Annie Beacon	3	282-802-5720
5	Marjorie Waller	1	807-221-8827
6	John Chan	5	727-315-7294
7	Satish Dhawan	2	294-849-1729
8	Wilhelm Steinberg	1	749-998-1556
9	Total	=sum(b2:b8)	




Modest impact

Massive impact

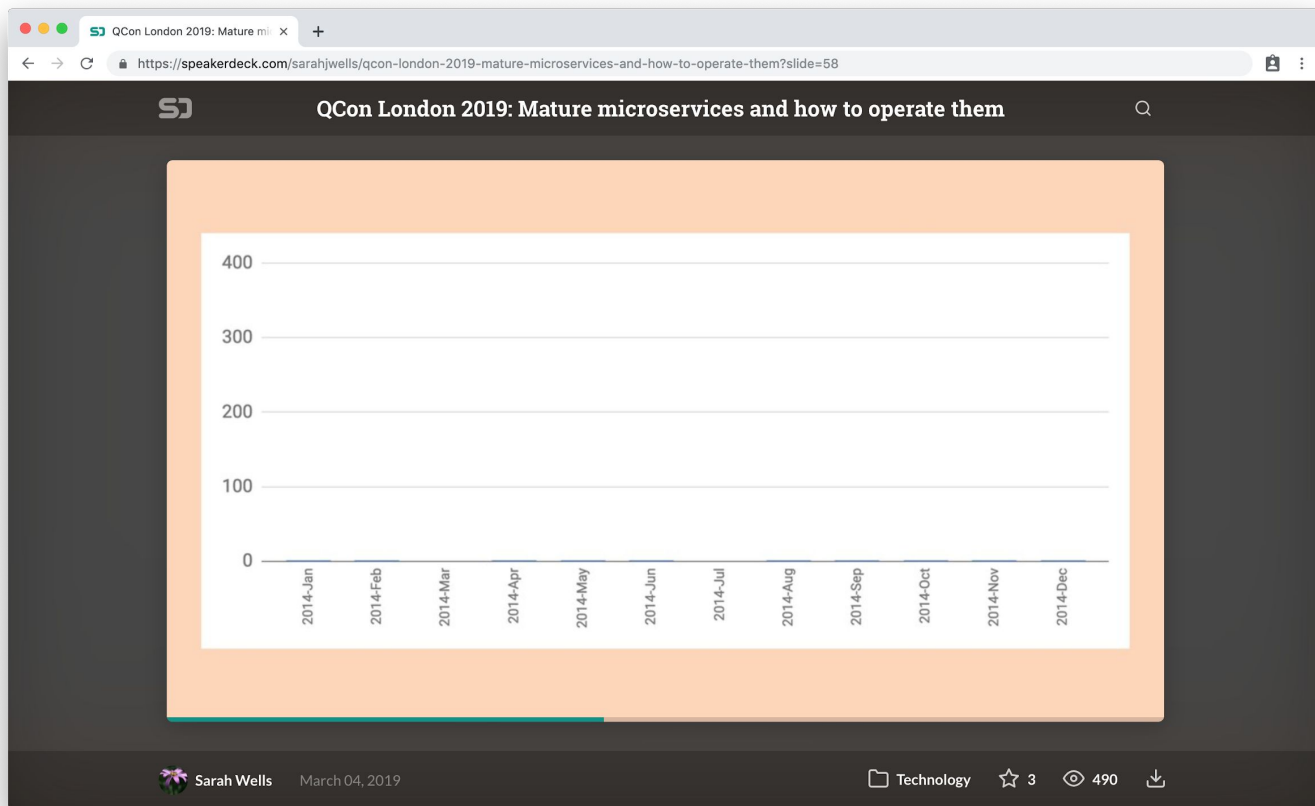
Small software scope

Large software scope

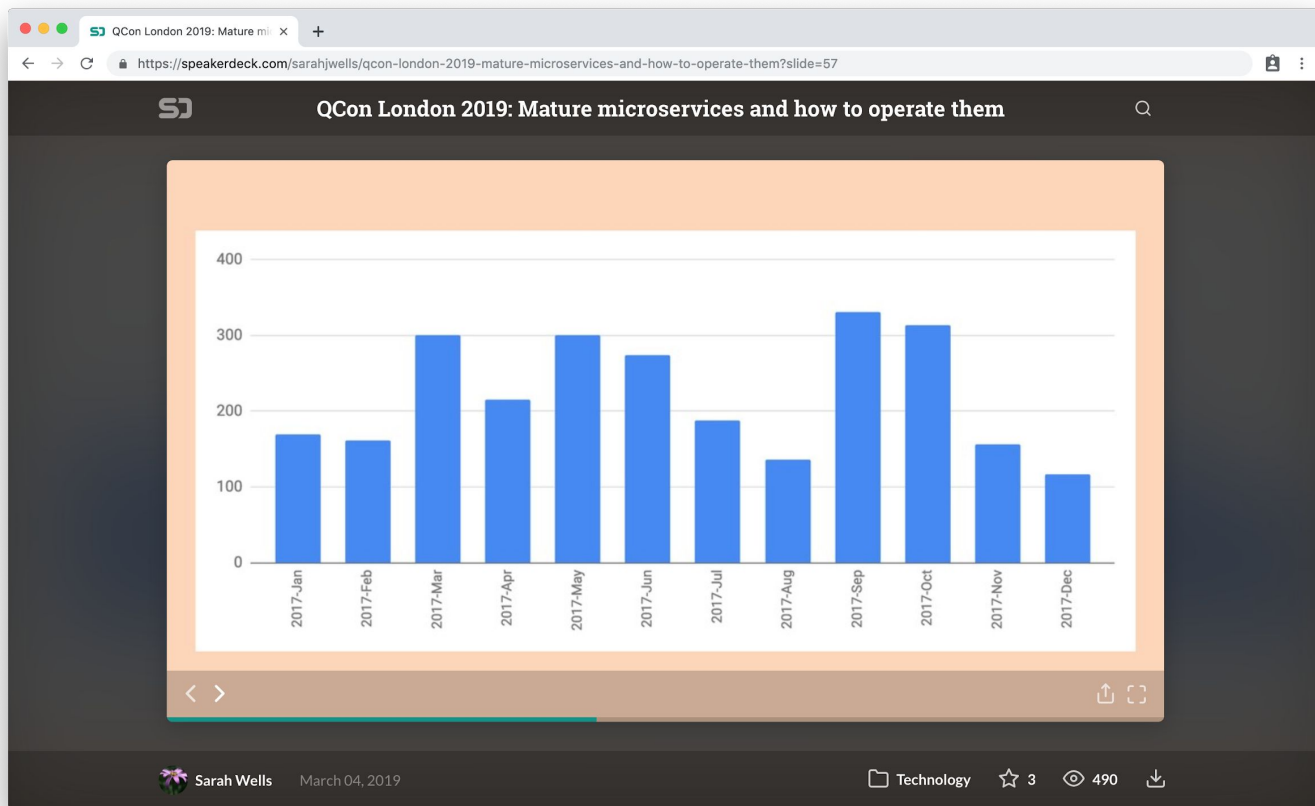
Developers: 



@sarahwells: monthly release count with a monolith




@sarahwells: monthly release count with microservices



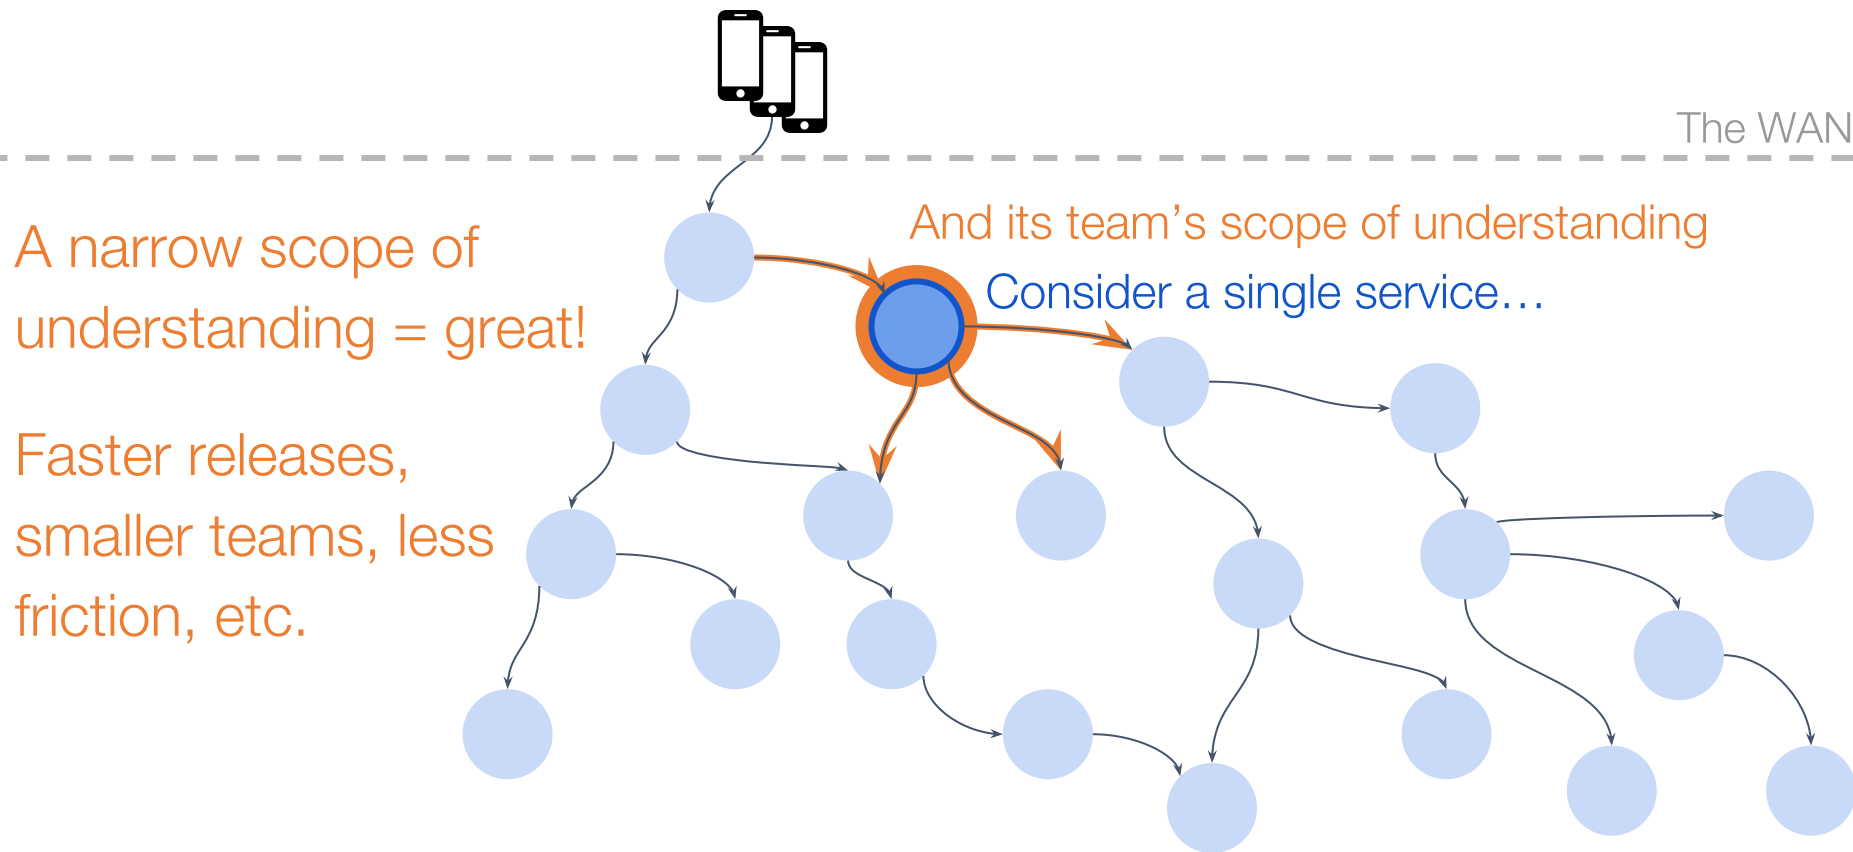
A microservices architecture



The WAN



A microservices architecture




A microservices architecture (with a slowdown)



The WAN

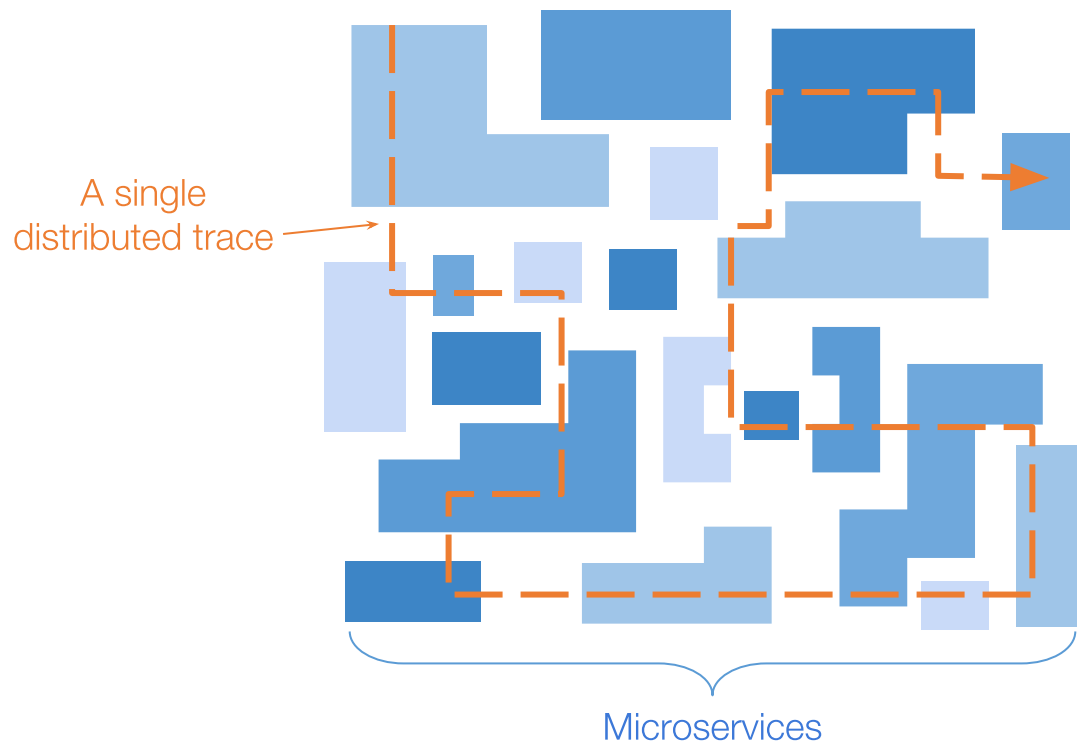
A narrow scope of
understanding = great!
... problematic.

Decoupling cuts
both ways.




Crisis of Confidence #2

Hypothesis Overload



Tracing Conundrum: Data Volume



app transaction rate
x # of microservices
x cost of net+storage
x weeks of retention

way too much \$\$\$\$

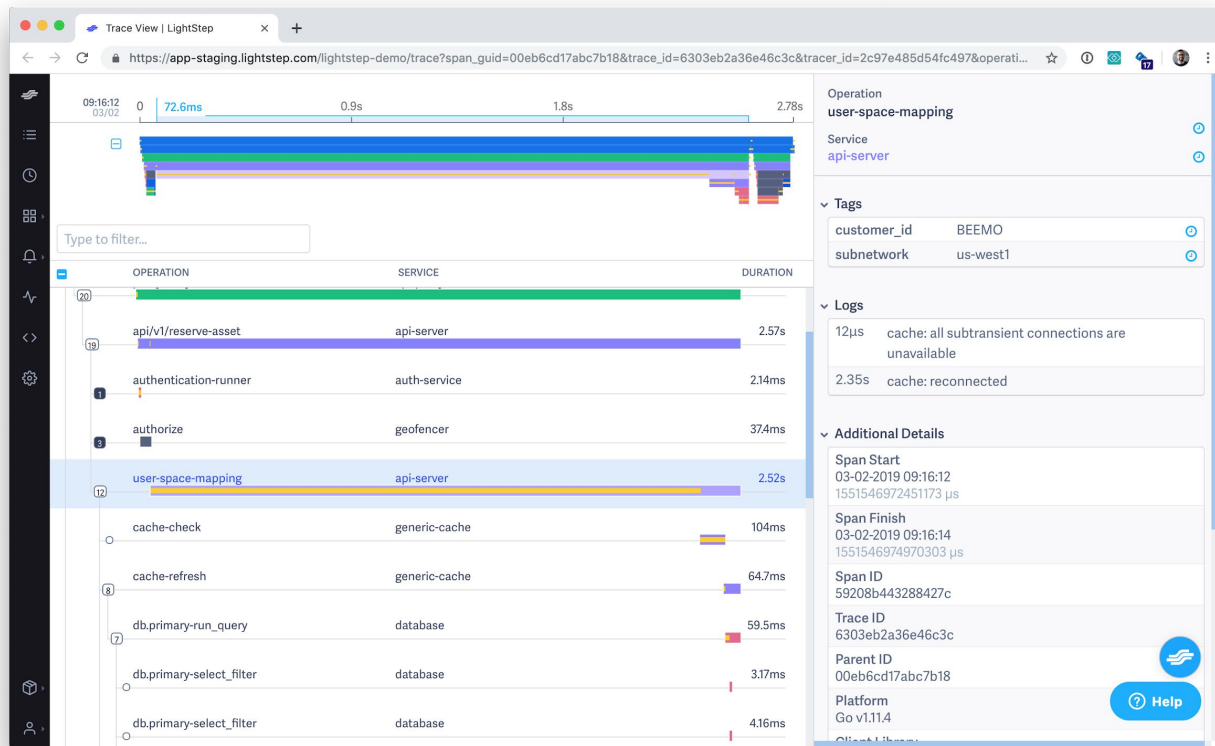
The Life of Trace Data: Dapper

Stage	Overhead affects...	Retained
Instrumentation Executed	App	100.00%
Buffered within app process	App	000.10%
Flushed out of process	App	000.10%
Centralized regionally	Regional network + storage	000.10%
Centralized globally	WAN + storage	000.01%

The Life of Trace Data: ~~Dapper~~ “Other Approaches”

Stage	Overhead affects...	Retained
Instrumentation Executed	App	100.00%
Buffered within app process	App	100.00%
Flushed out of process	App	100.00%
Centralized regionally	Regional network + storage	100.00%
Centralized globally	WAN + storage	on-demand

Tracing Conundrum: “Too Much Signal”



Raw distributed traces, in summary

- One distributed trace per transaction
- Crosses microservice boundaries
- They are **necessary** if we want to understand the relationships between distant actors in our architecture ...

... and yet:

- traces are too big for our brains
- there are too many of them to centralize durably

“Traces” vs. “Tracing”

Distributed *traces*: basically just structs

Distributed *tracing*: **the art and science of making distributed traces valuable**

So... how *do* we make
distributed traces valuable?

We use them to restore confidence.

Traces, Tracing, and Confidence

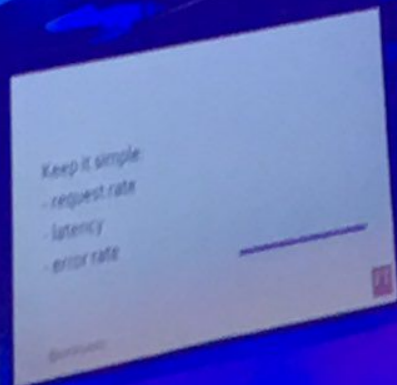
Quick Vocab Refresher: **SLIs**

“SLI” = “Service Level Indicator”

TL;DR: An SLI is **an indicator of health** that a service’s **consumers** would care about.

... *not* an indicator of its inner workings

A blast from the past...



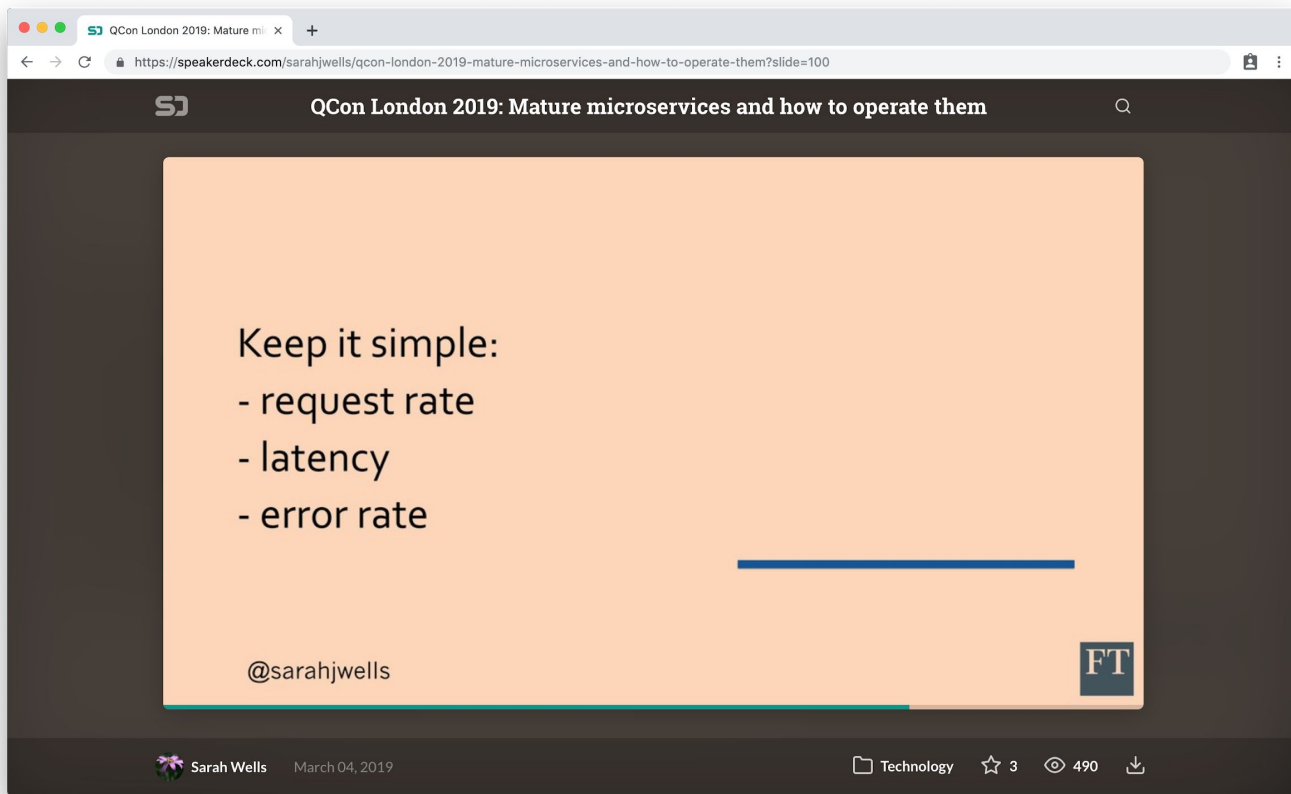
SLI advice from earlier to now

QCon

QCon

QCon

Sarah's actual slide, about SLIs



The image is a screenshot of a web browser displaying a presentation slide. The browser's address bar shows the URL <https://speakerdeck.com/sarahjwells/qcon-london-2019-mature-microservices-and-how-to-operate-them?slide=100>. The slide itself has a dark header with the text "QCon London 2019: Mature microservices and how to operate them" and a search icon. The main content area is a light orange rectangle with the text "Keep it simple:" followed by a bulleted list: "- request rate", "- latency", and "- error rate". Below the list is a blue horizontal line. At the bottom left of the slide is the text "@sarahjwells" and at the bottom right is the "FT" logo. The browser's footer shows the speaker's name "Sarah Wells", the date "March 04, 2019", and engagement metrics: "Technology", "3" stars, "490" views, and a download icon.

QCon London 2019: Mature microservices and how to operate them

Keep it simple:

- request rate
- latency
- error rate

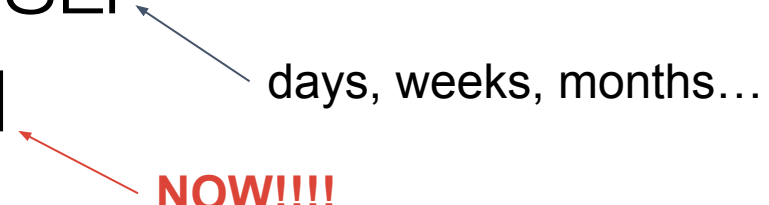
@sarahjwells

FT

Sarah Wells March 04, 2019 Technology 3 490

“Confidence” is controlling your SLIs

Two fundamental goals:

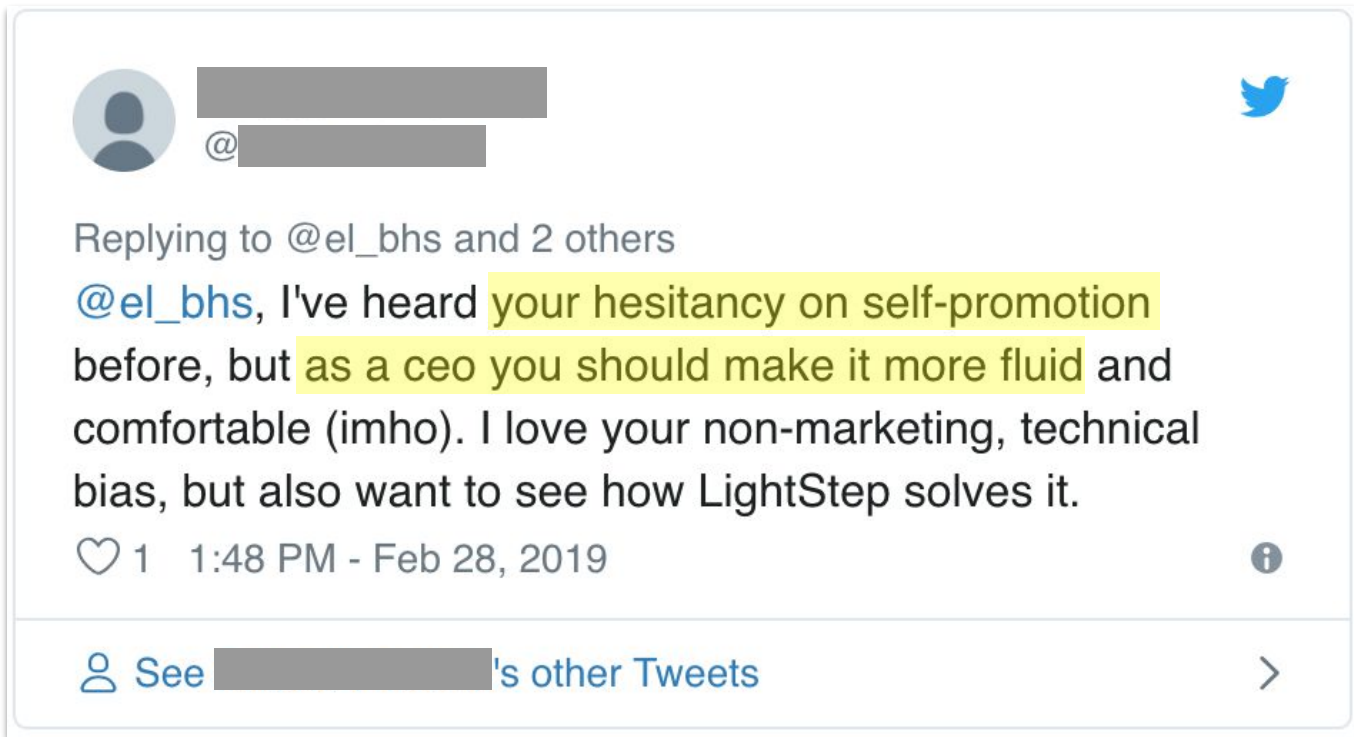
- Gradually **improving** an SLI
 - Rapidly **restoring** an SLI
- 
- days, weeks, months...
- NOW!!!!**

Reminder: “SLI” = “Service Level Indicator”

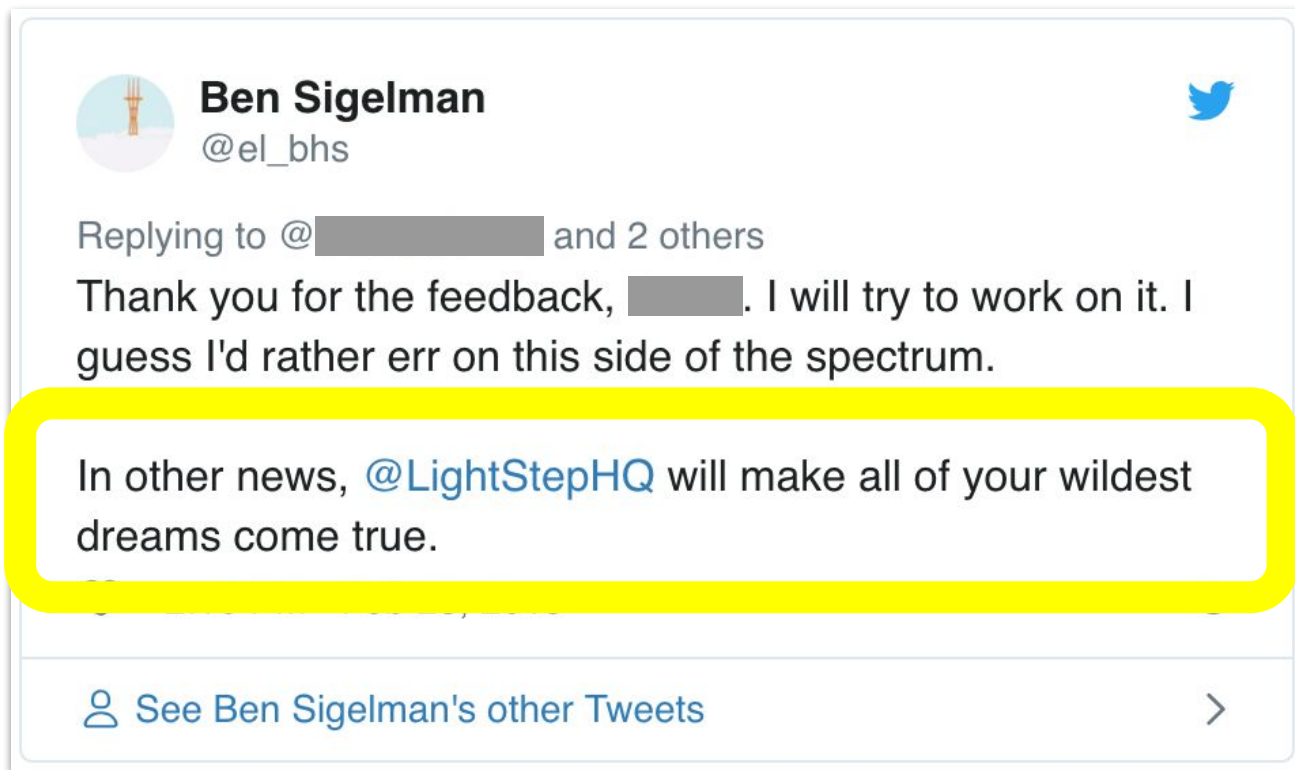
Two Fundamental **Activities**

1. **Detection:** measuring SLIs precisely
2. **Explaining variance:** recognizing and explaining variance, often iteratively

An aside...



An aside...




A Service-Centric Approach



The WAN

Given any service:

1. Start with an SLI
2. Find variance
3. Explain it

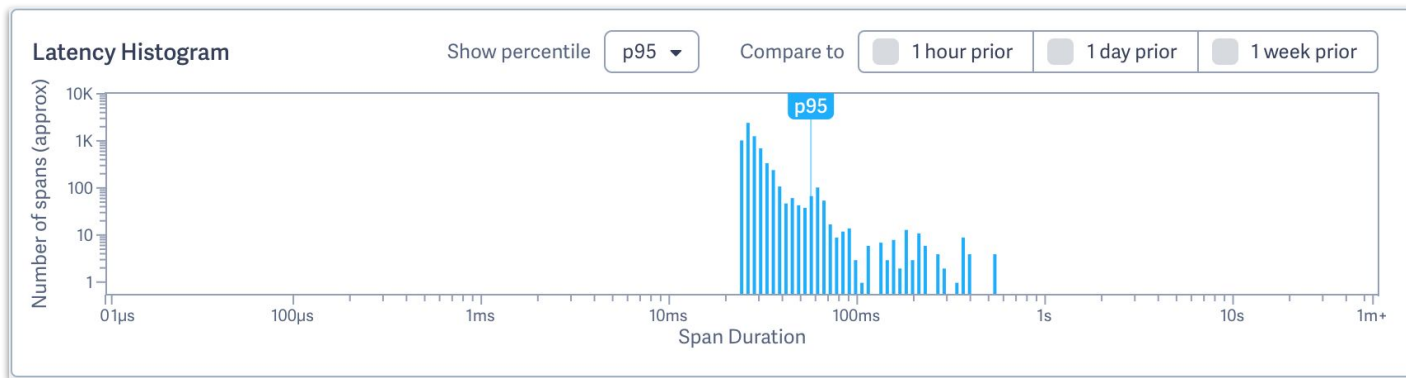


Confidence about *variance*


Percentile	Latency for A	Latency for B
p50	30ms	3ms
p90	42ms	11ms
p95	66ms	13ms
p99	180ms	17ms

Confidence about *variance*


A




B



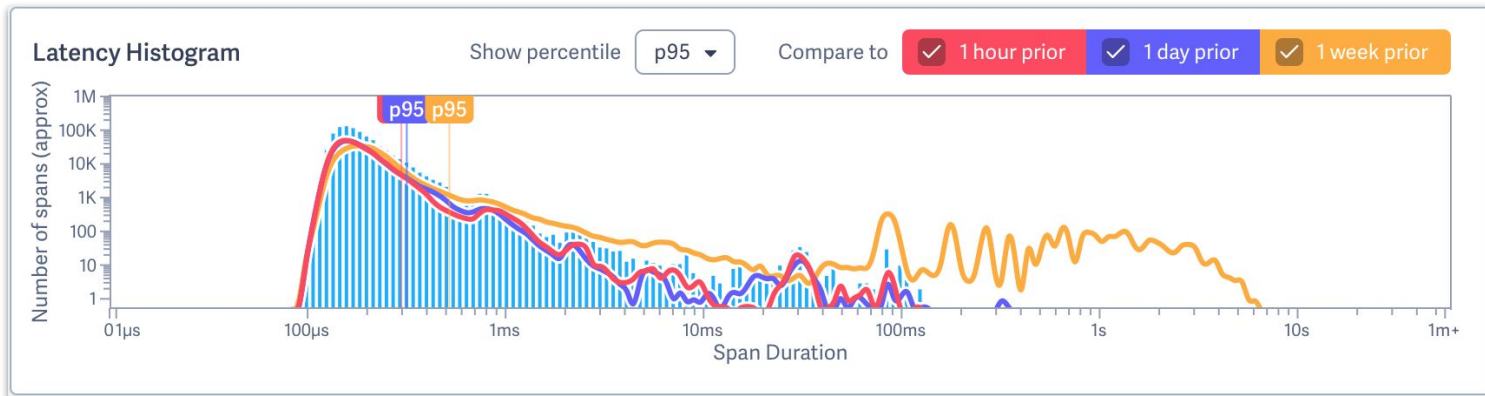

Confidence about *variance*



Confidence about *variance*



Latency variance vs temporal variance

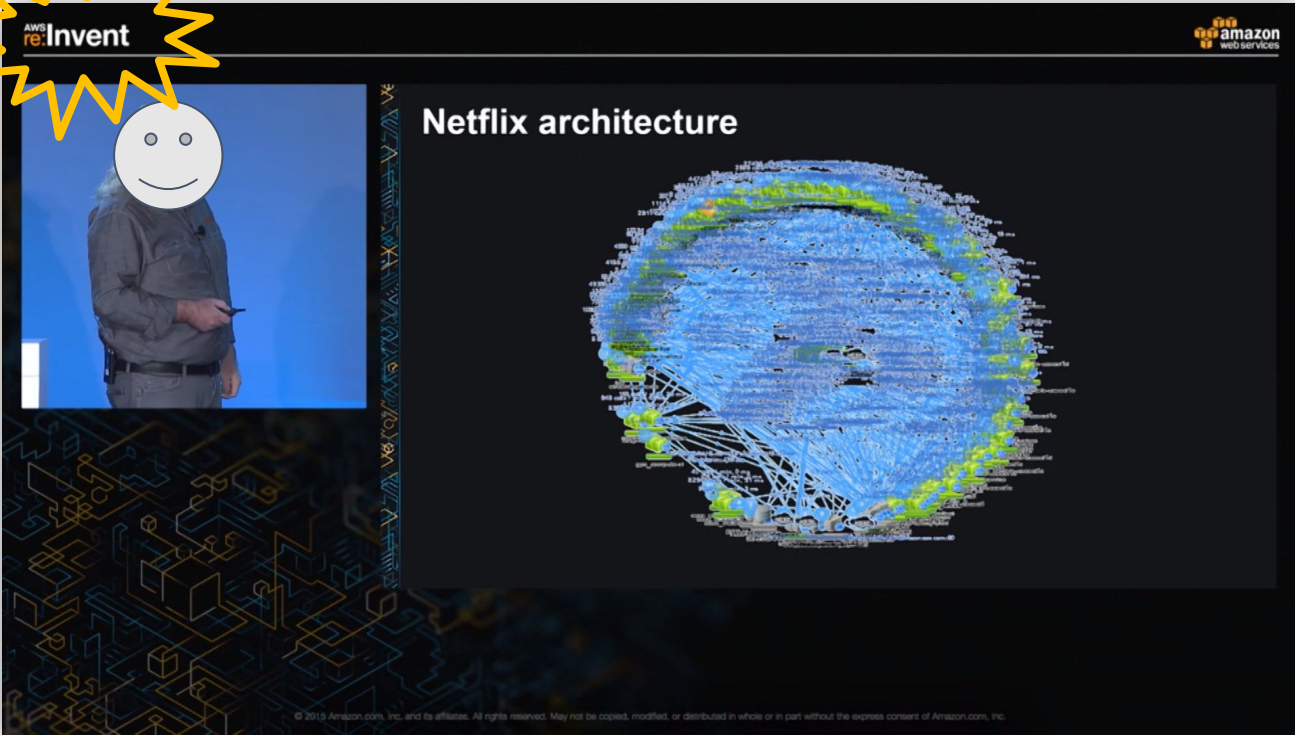


“Variance over *what* exactly?”

just trace span latencies,

filtered by (high-cardinality) tag values

Confidence about *bottlenecks*

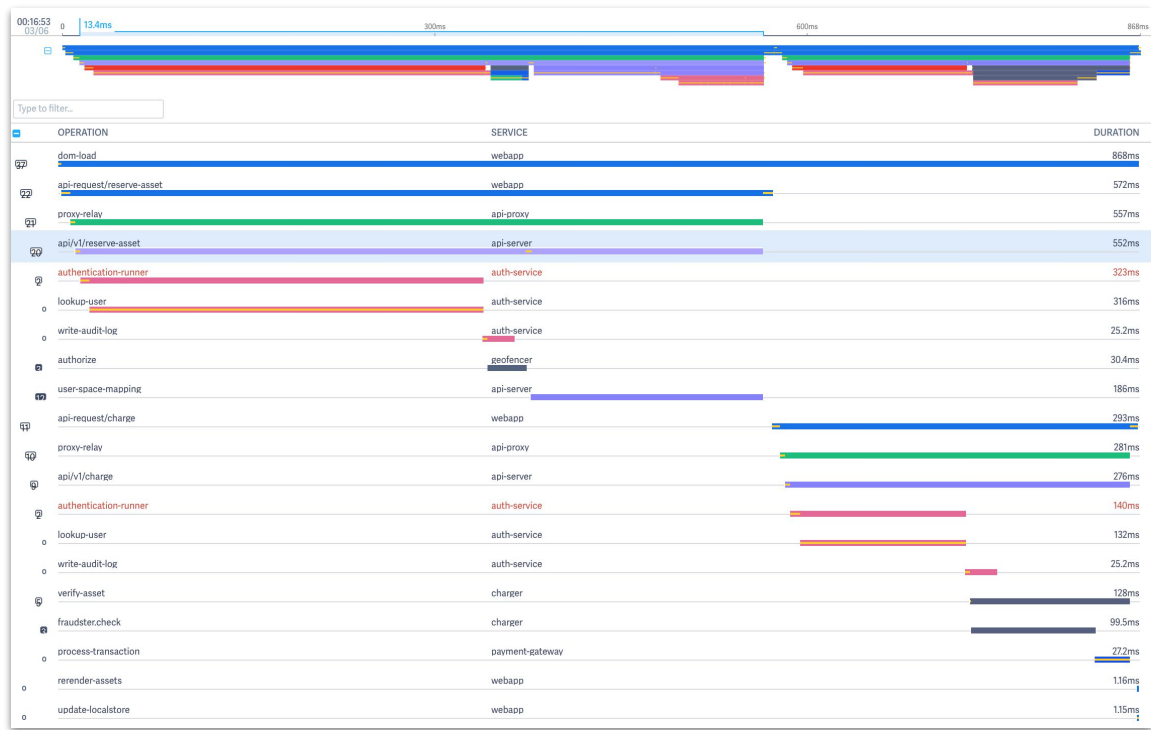


The image is a screenshot of a presentation slide from AWS re:Invent. The slide features a dark background with a decorative pattern of circuit-like lines in blue and yellow. On the left side, there is a photograph of a man in a grey jacket, with a white smiley face icon placed over his head. Above the photo is a yellow jagged starburst graphic. The top left corner of the slide has the 'AWS re:Invent' logo, and the top right corner has the 'amazon web services' logo. The main title of the slide is 'Netflix architecture' in white text. Below the title is a large, complex network diagram representing the Netflix architecture. The diagram consists of numerous nodes, some colored blue and others green, connected by a dense web of white lines. The nodes are arranged in a roughly circular pattern, with some lines extending outwards. At the bottom of the slide, there is a small copyright notice: '© 2015 Amazon.com, Inc. and its affiliates. All rights reserved. May not be copied, modified, or distributed in whole or in part without the express consent of Amazon.com, Inc.'

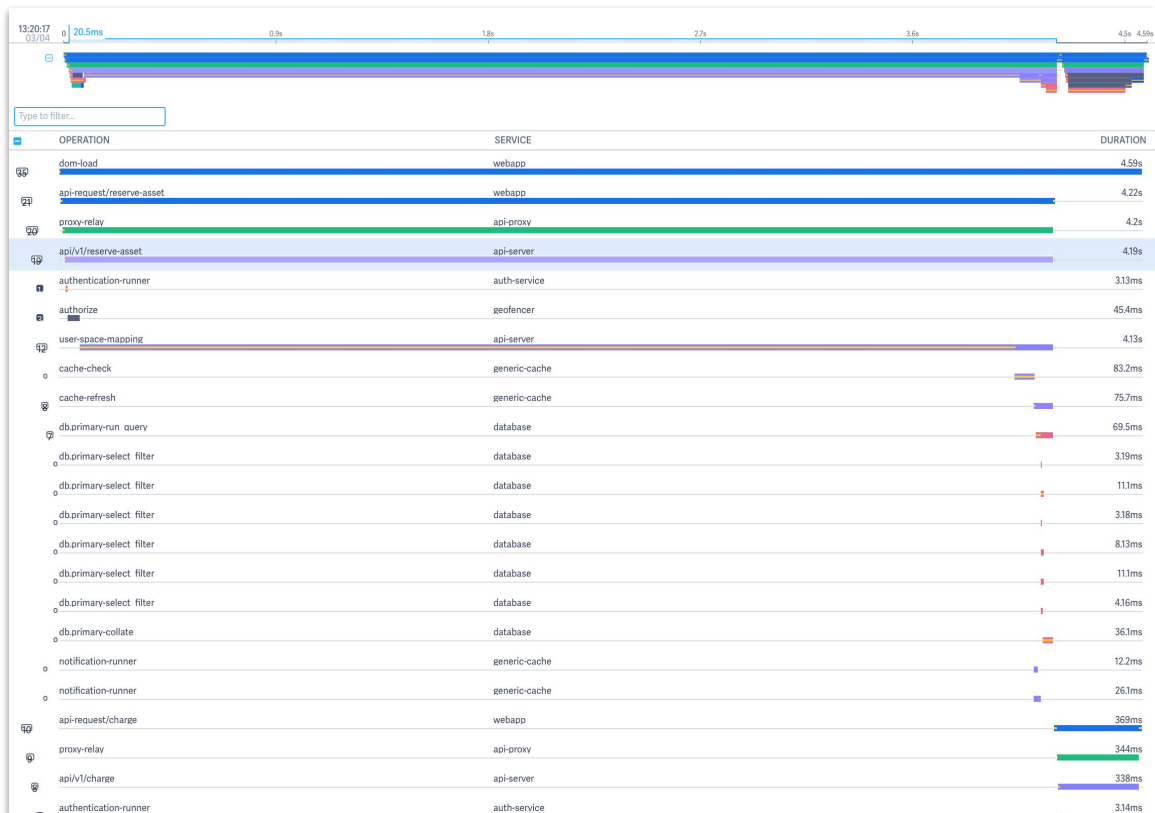
Confidence about *bottlenecks*

([link](#))

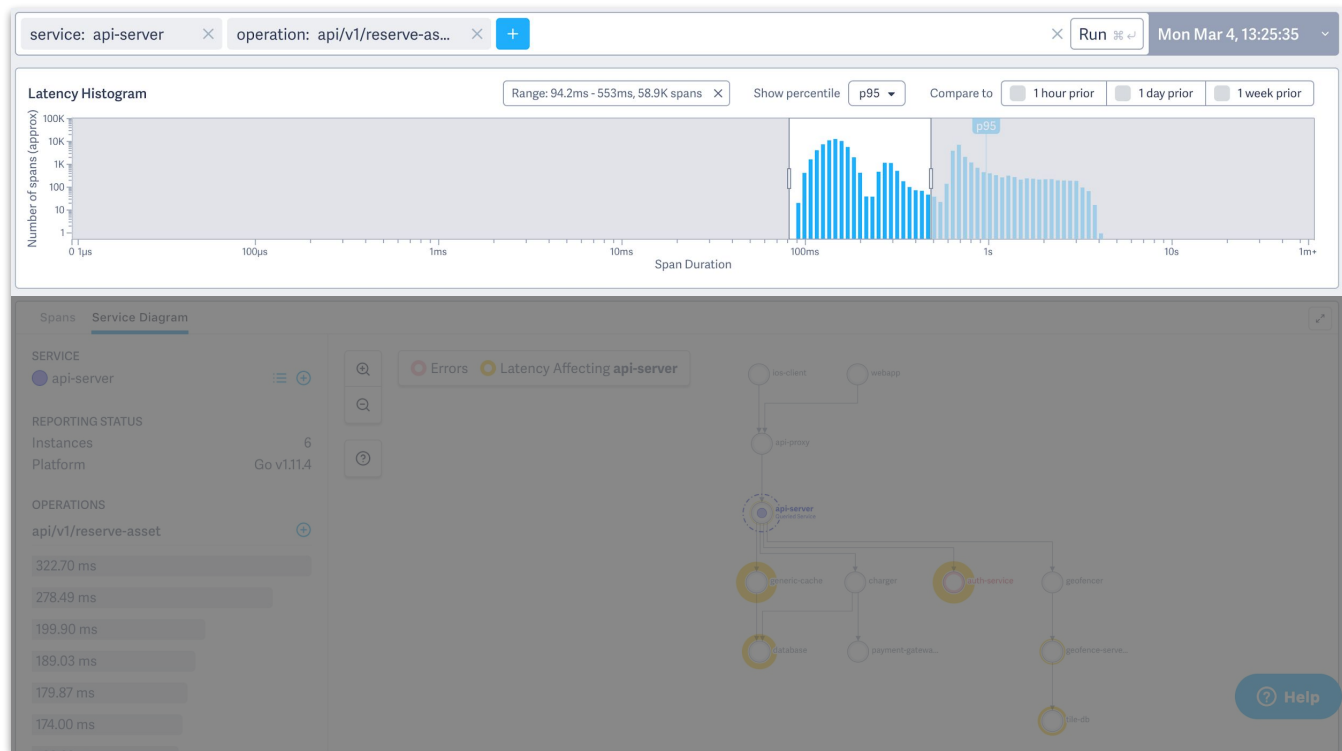
Common case traces



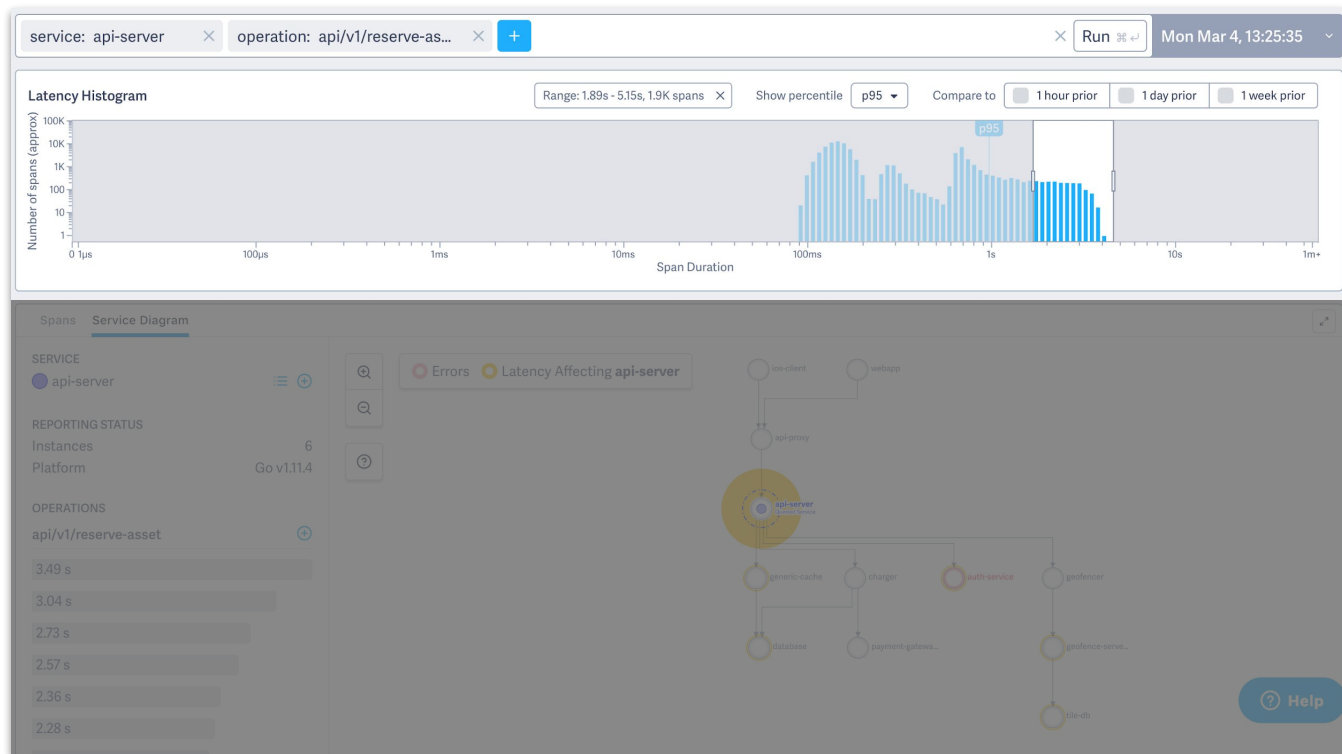
Common case traces



Confidence about *bottlenecks*: common case



Confidence about *bottlenecks*: outliers



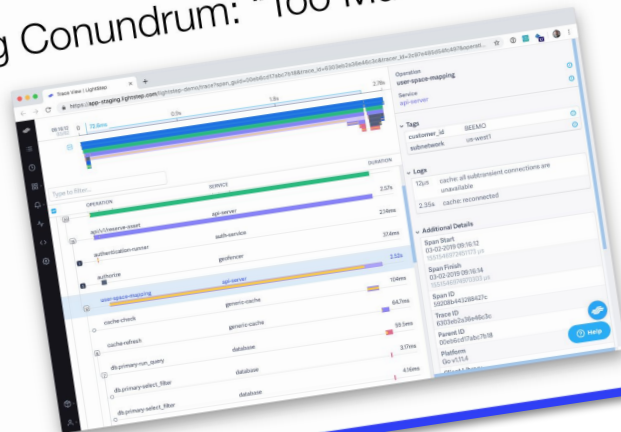
Confidence about *hypotheses*

(link)

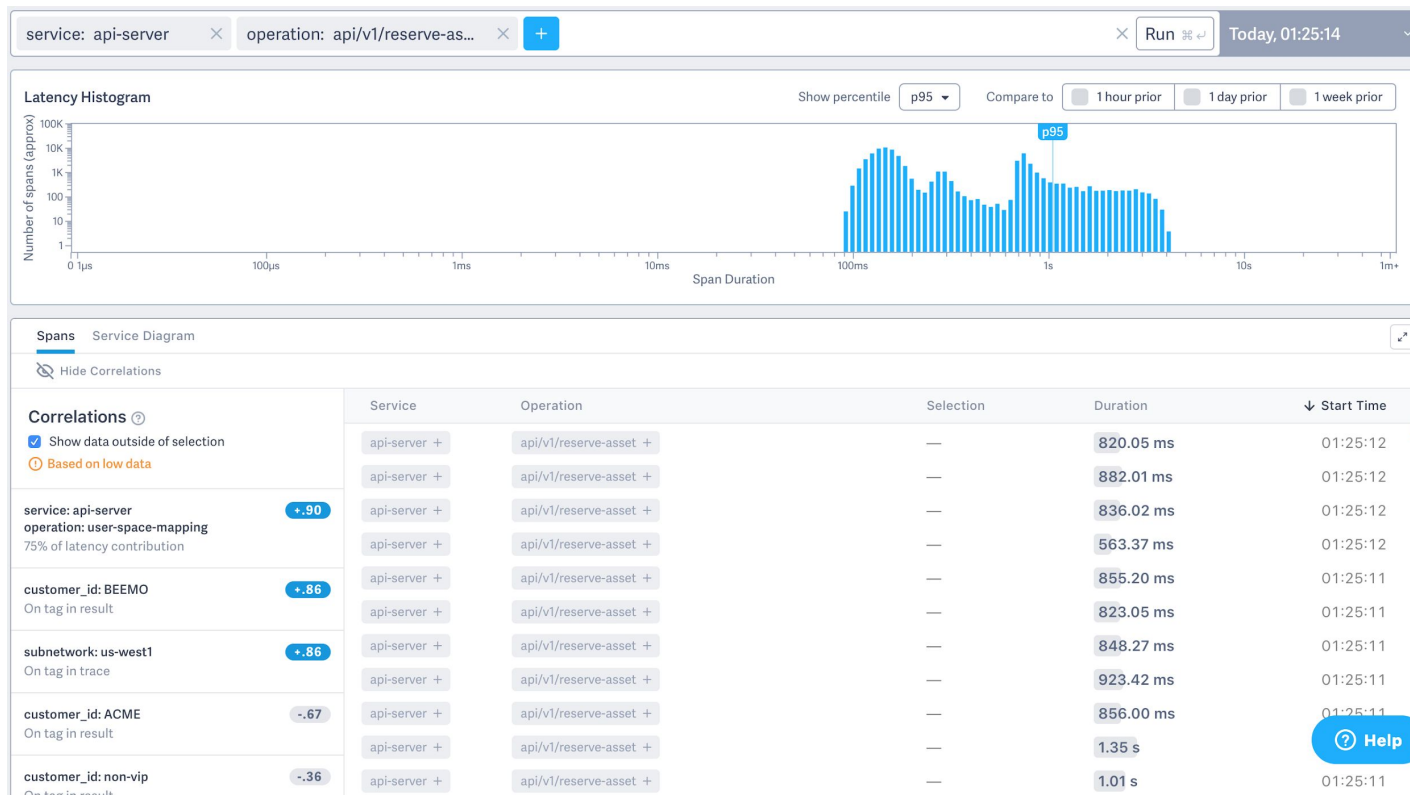
(Relevant)



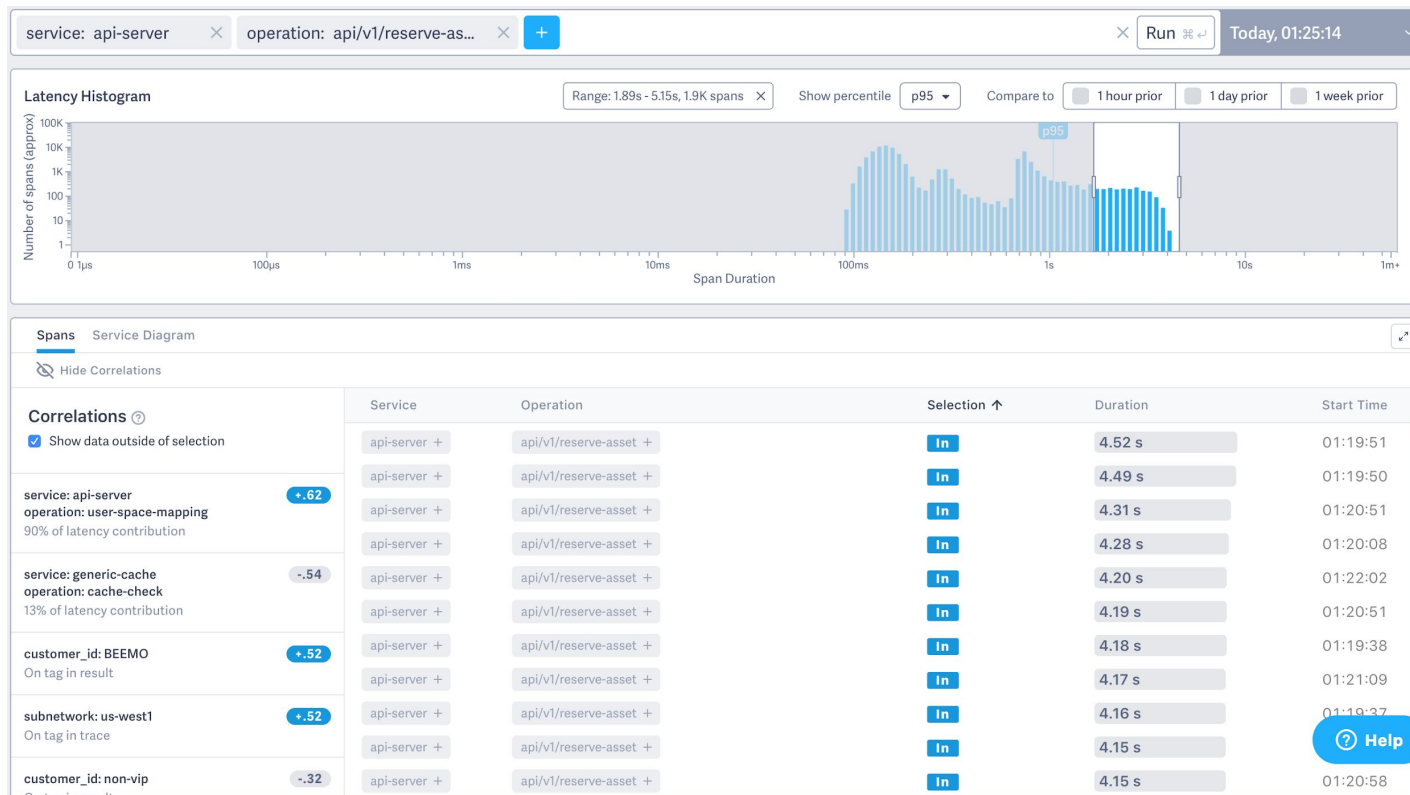
Tracing Conundrum: “Too Much Signal”



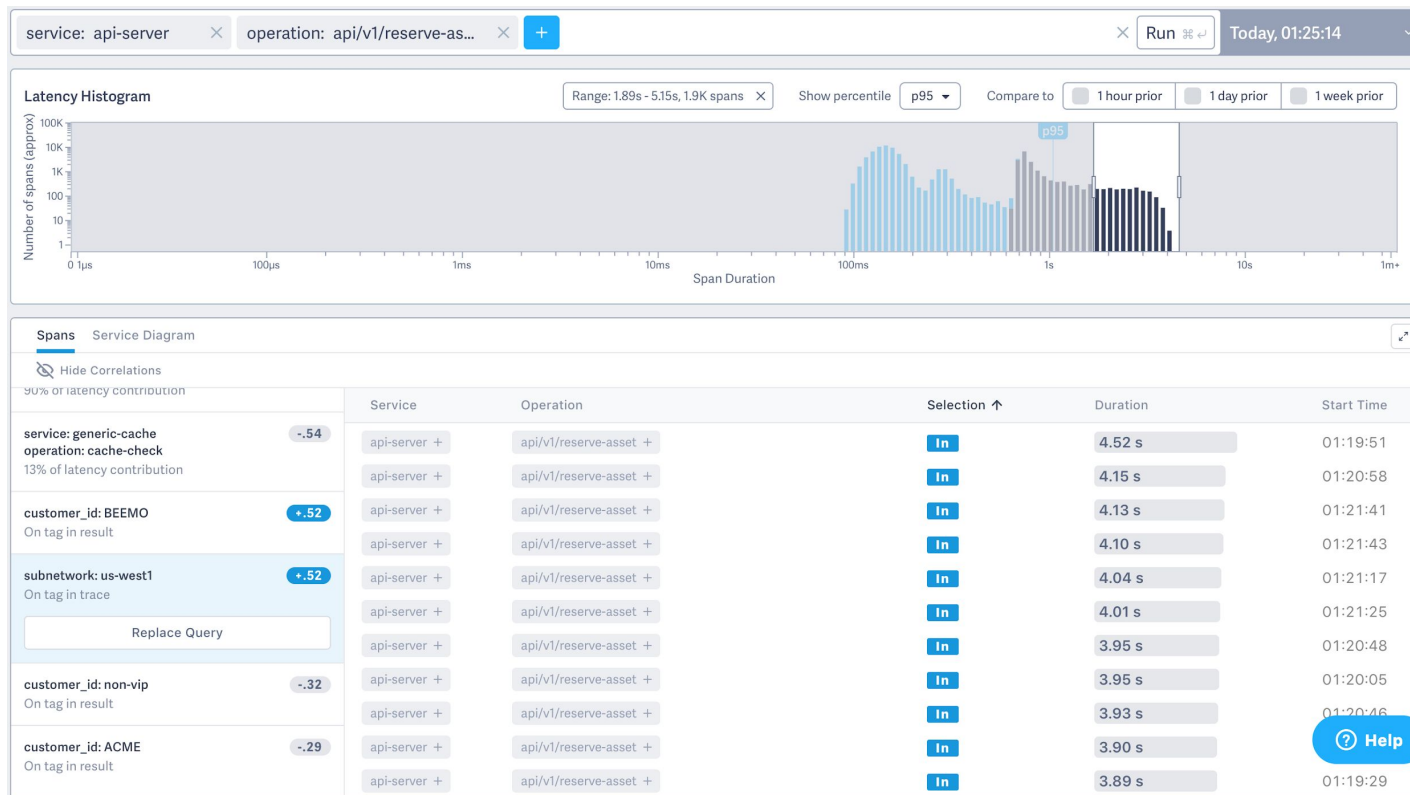
Confidence about *hypotheses*



Isolating variance



Considering a hypothesis



Wrapping up...



confidence

Thank you!

Ben Sigelman, Co-founder and CEO
twitter: [@el_bhs](#)
email: bhs@lightstep.com

PS: LightStep launched self-serve tracing
at QCon London!

I am friendly and would love to
chat... please say hello, I don't
make it to Europe often!

