


The present and



FUTURE

of

Serverless

OBSERVABILITY

hi, my name is Yan.



hi, my name is Yan.

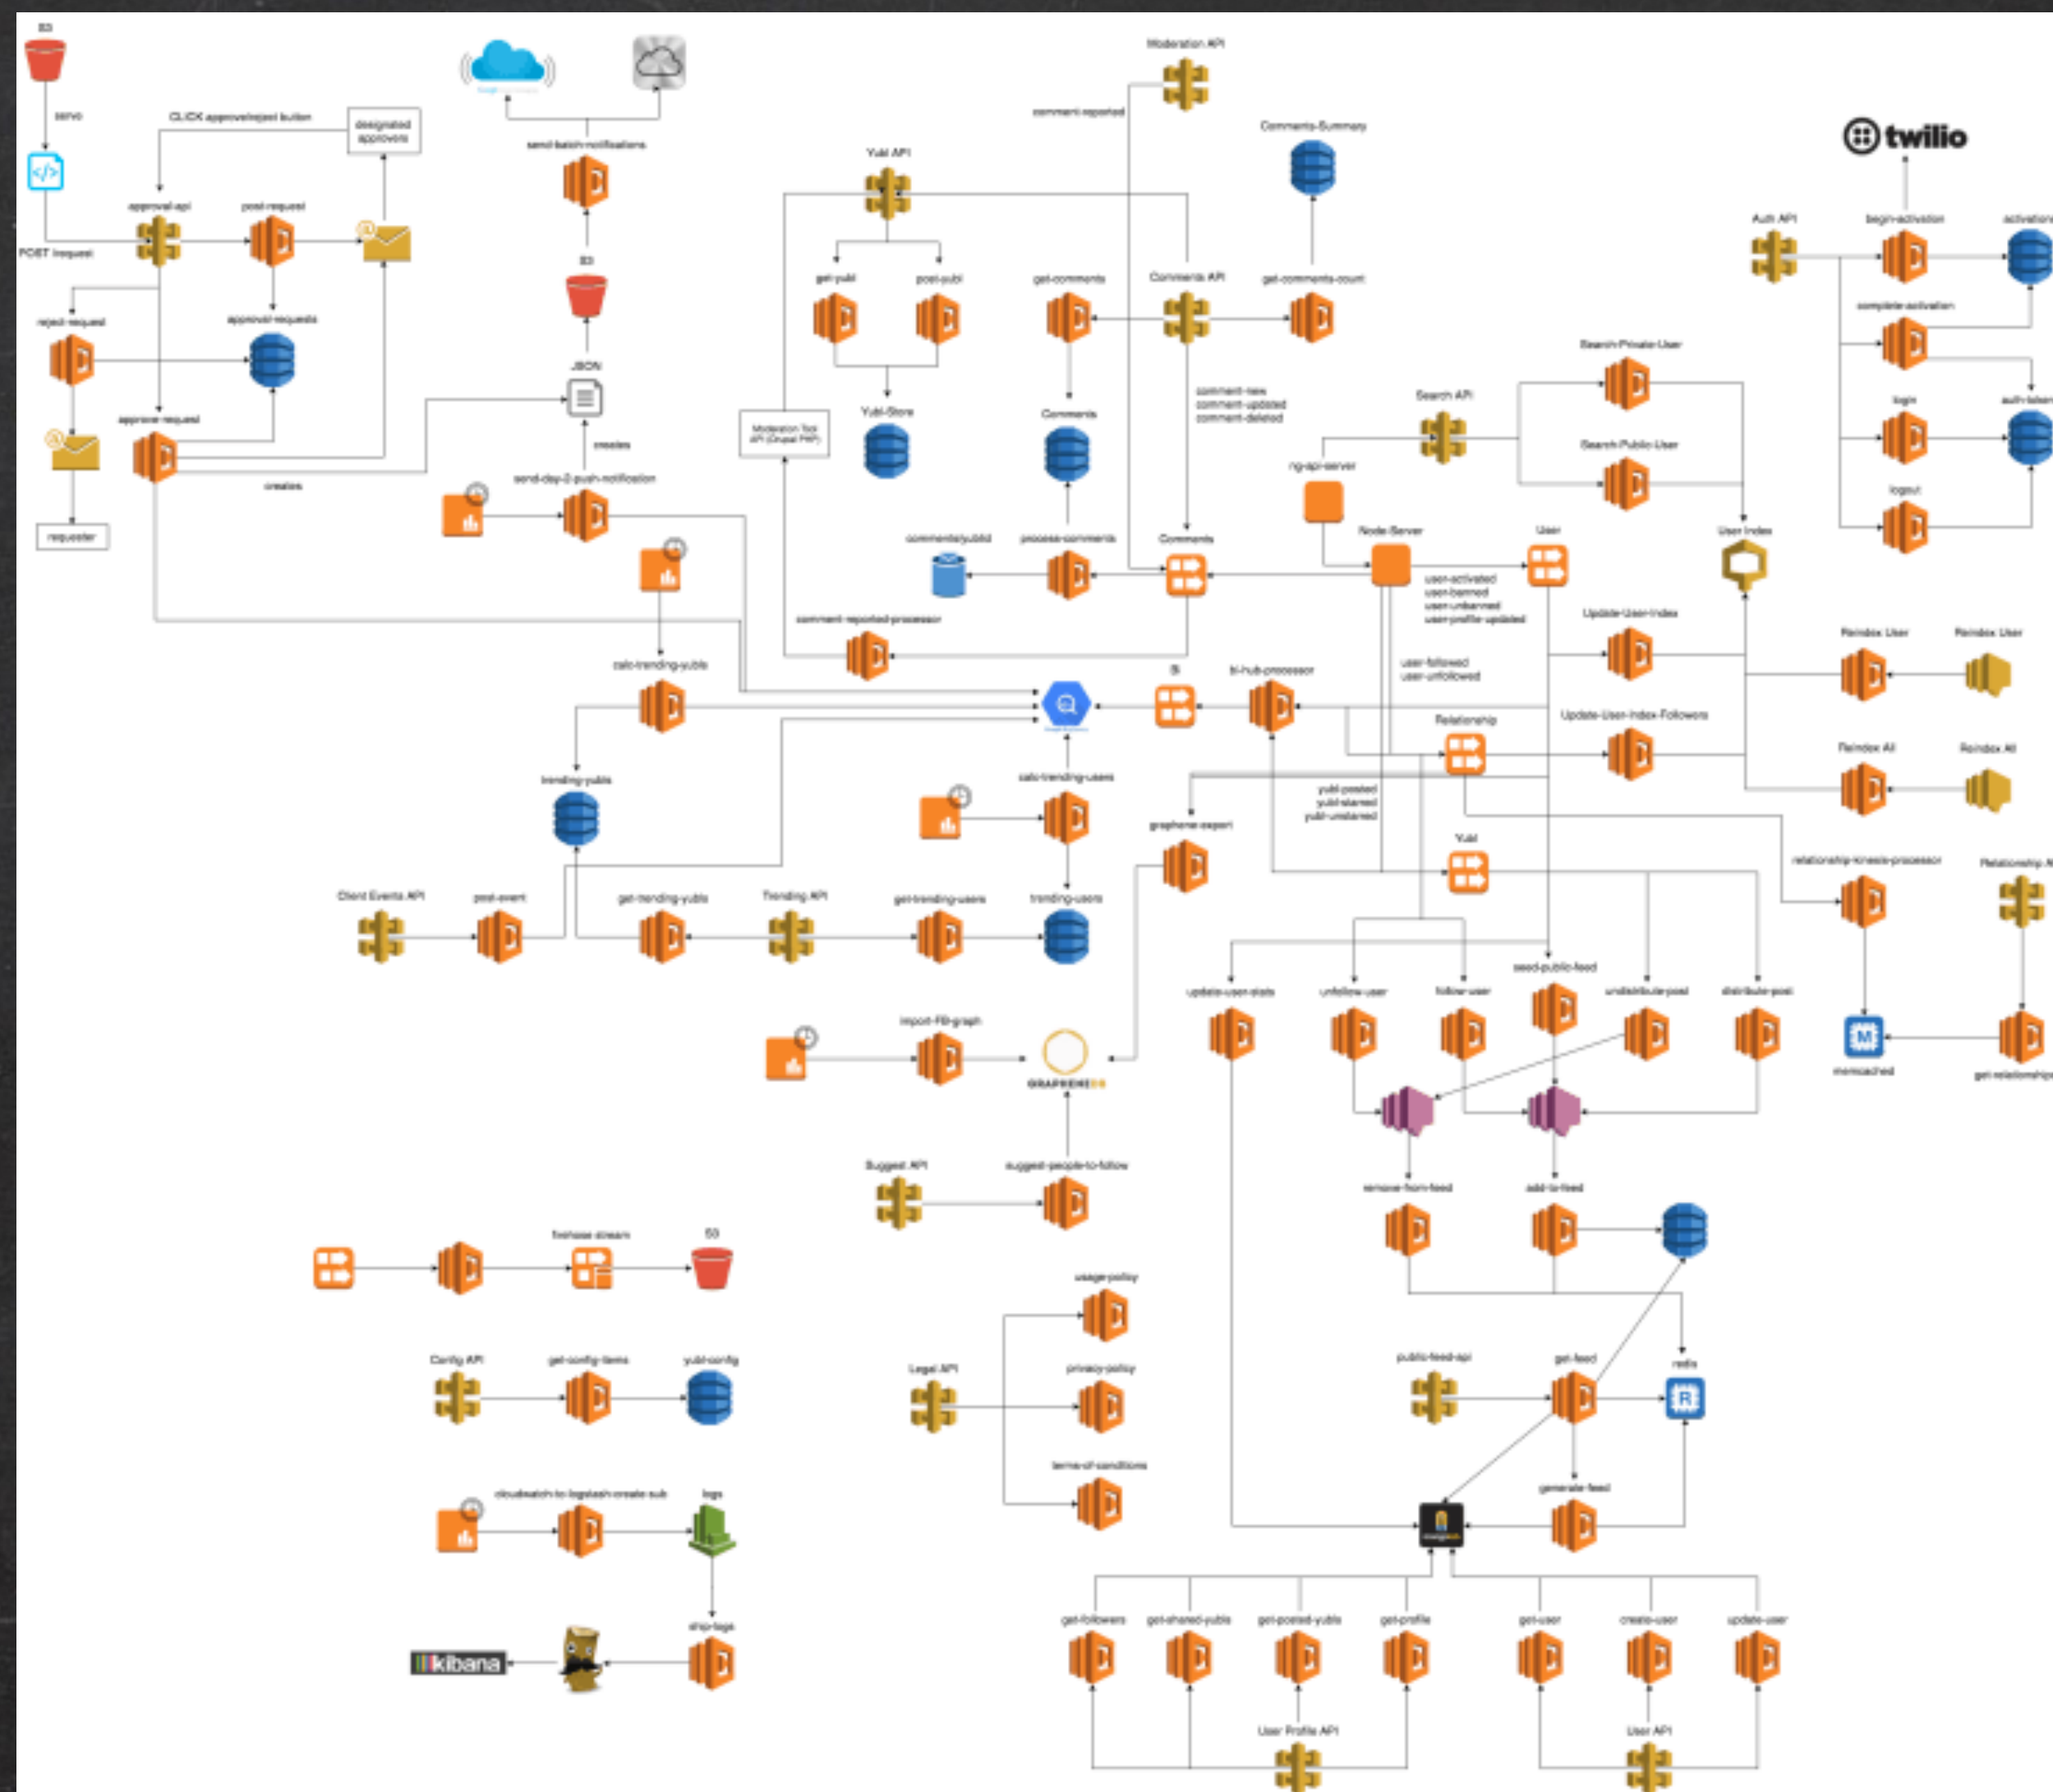


Space Ape™

hi, my name is Yan.



AWS user since 2009



<http://bit.ly/yubl-serverless>

Serverless Applications Lens

AWS Well-Architected Framework

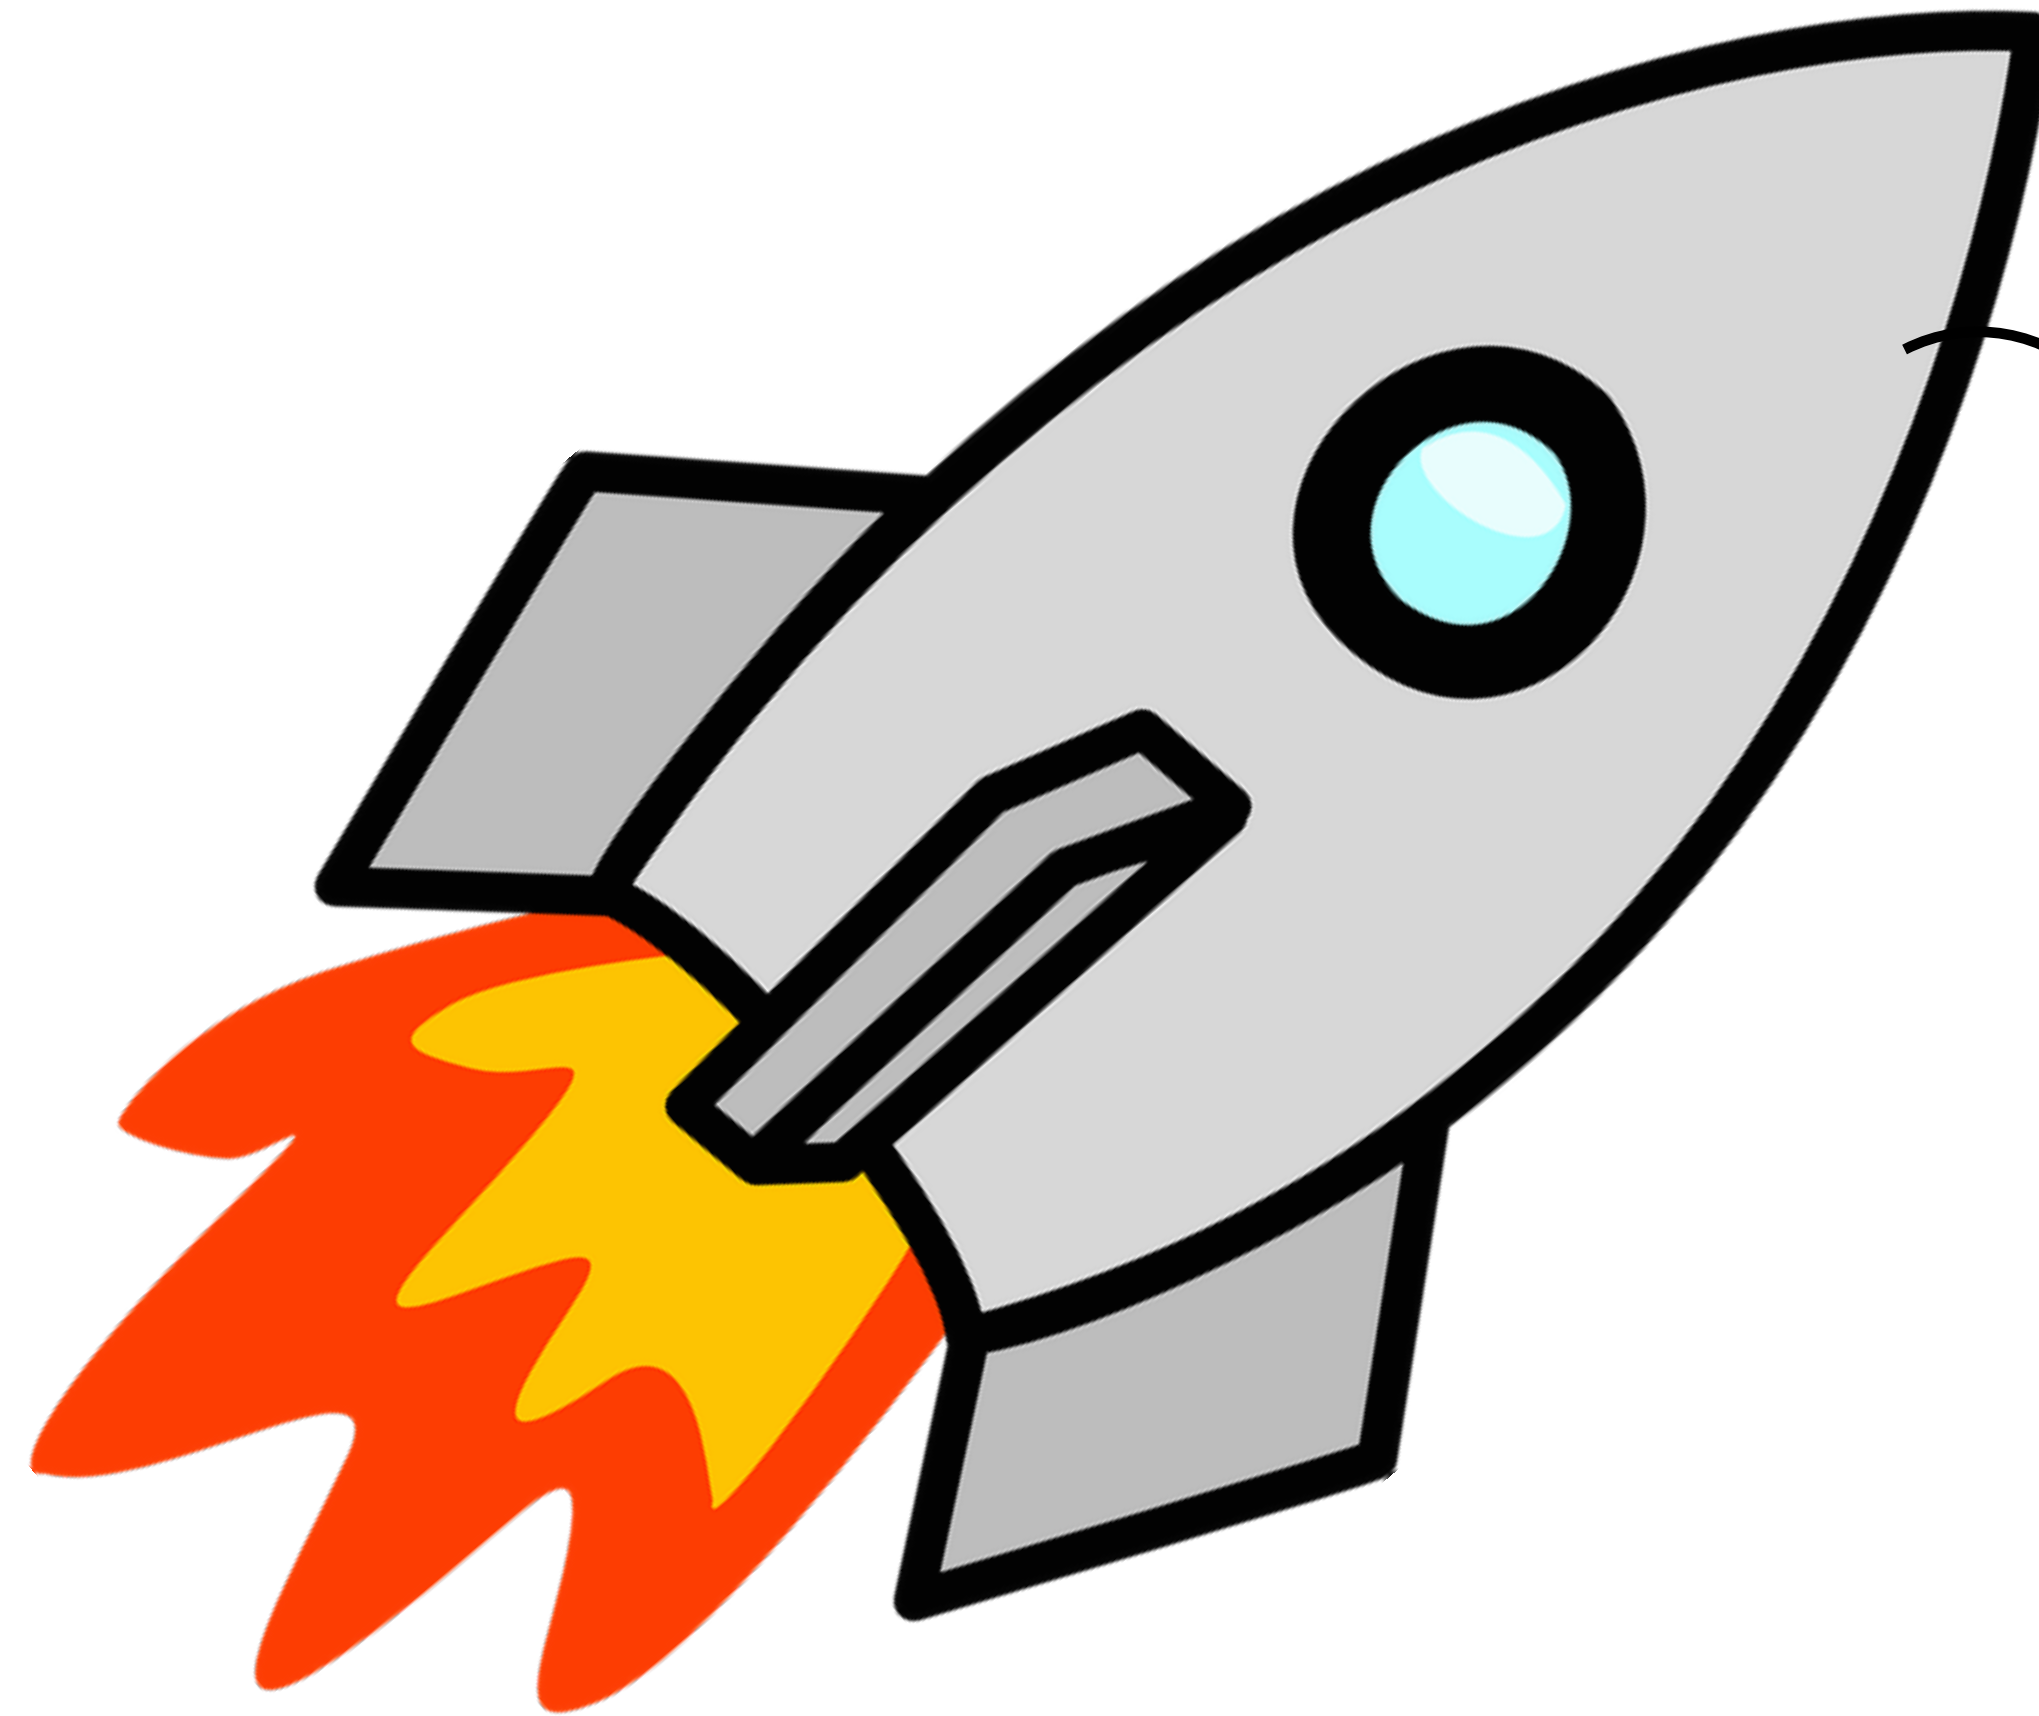
November 2017



<http://bit.ly/2Cdsai5>



2017



observability



Cindy Sridharan
@copyconstruct on Twitter
Sep 5, 2017 · 12 min read

Monitoring and Observability

During lunch with a few friends in late July, the topic of Observability came up. I have a talk coming up at Velocity in less than a month called ***Monitoring in the time of Cloud Native***, so I've been speaking with friends about how they approach monitoring where they work. During this conversation, one of my friends mentioned:



He was only half joking. I've heard several variations of this zinger, some of them being:

<http://bit.ly/2EXQZBj>

Observability at Twitter

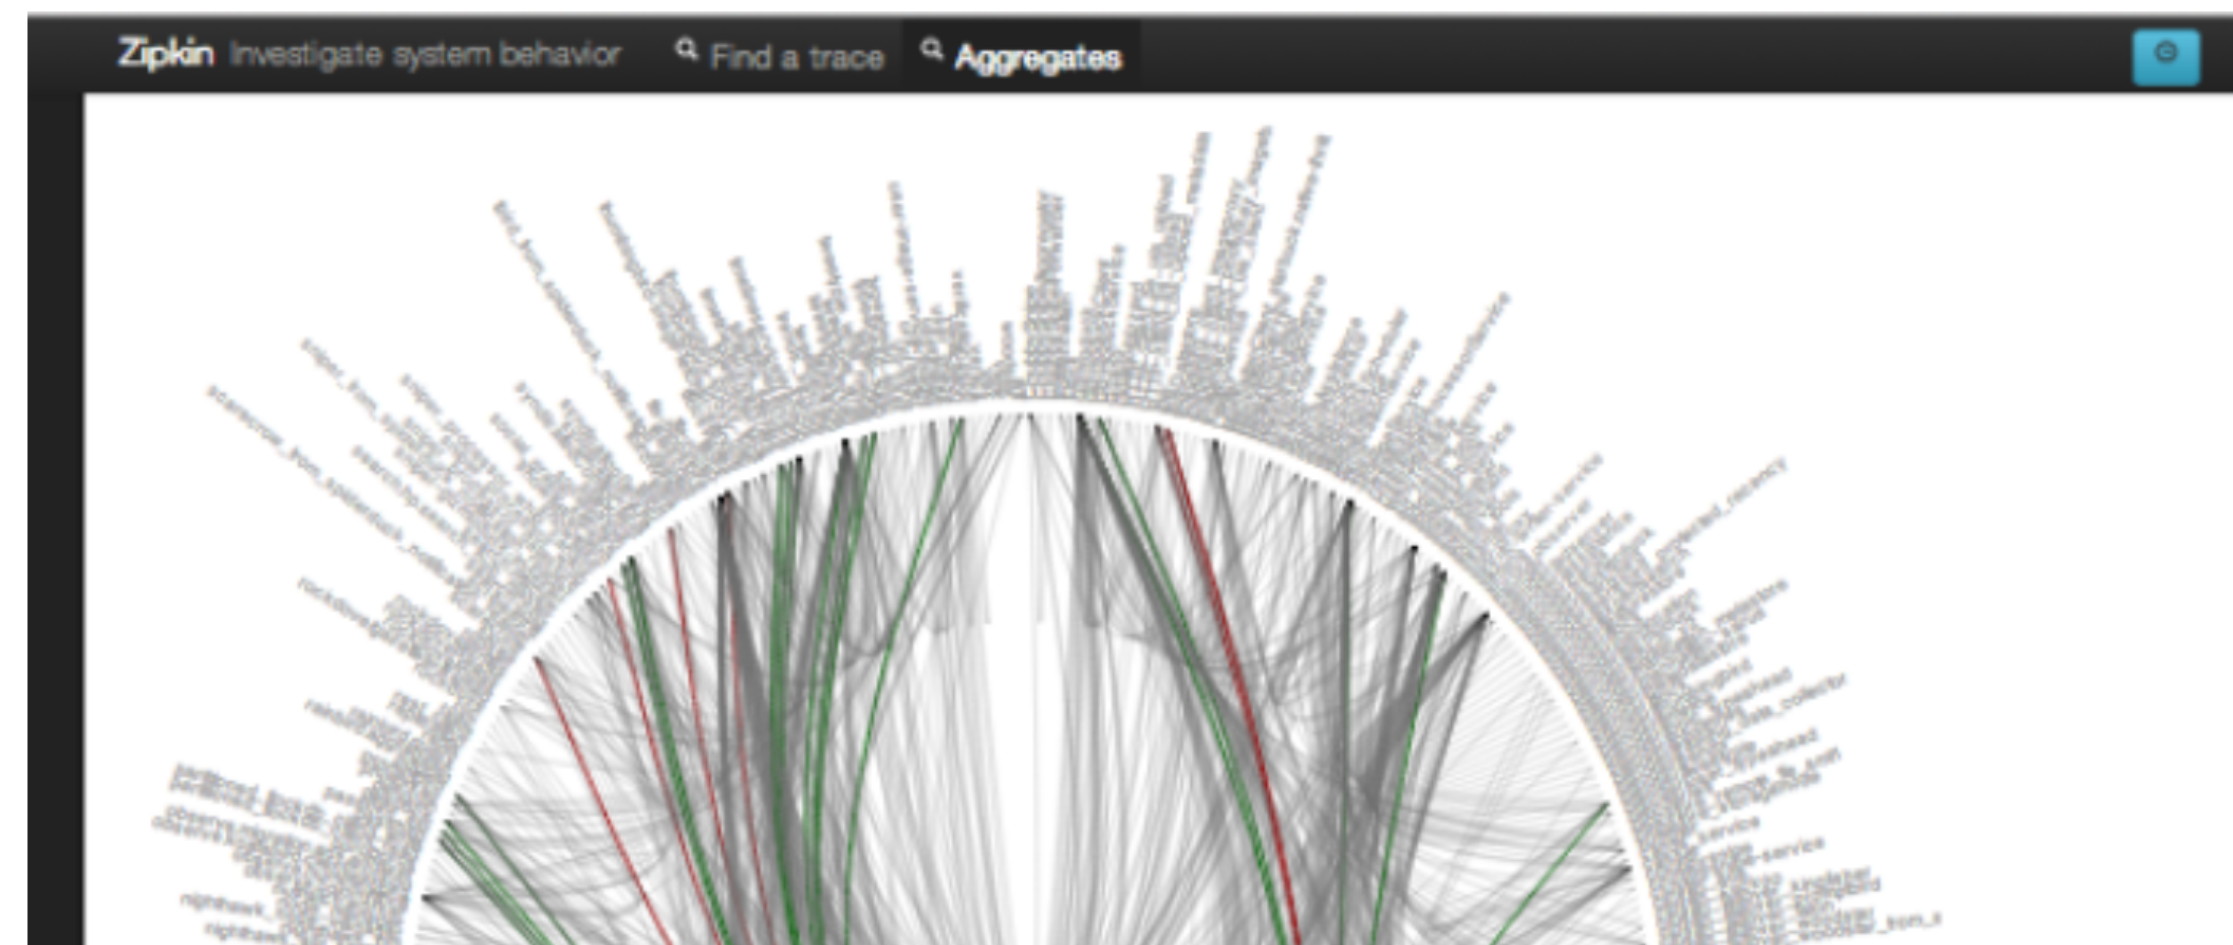
By @gphat

Monday, 9 September 2013

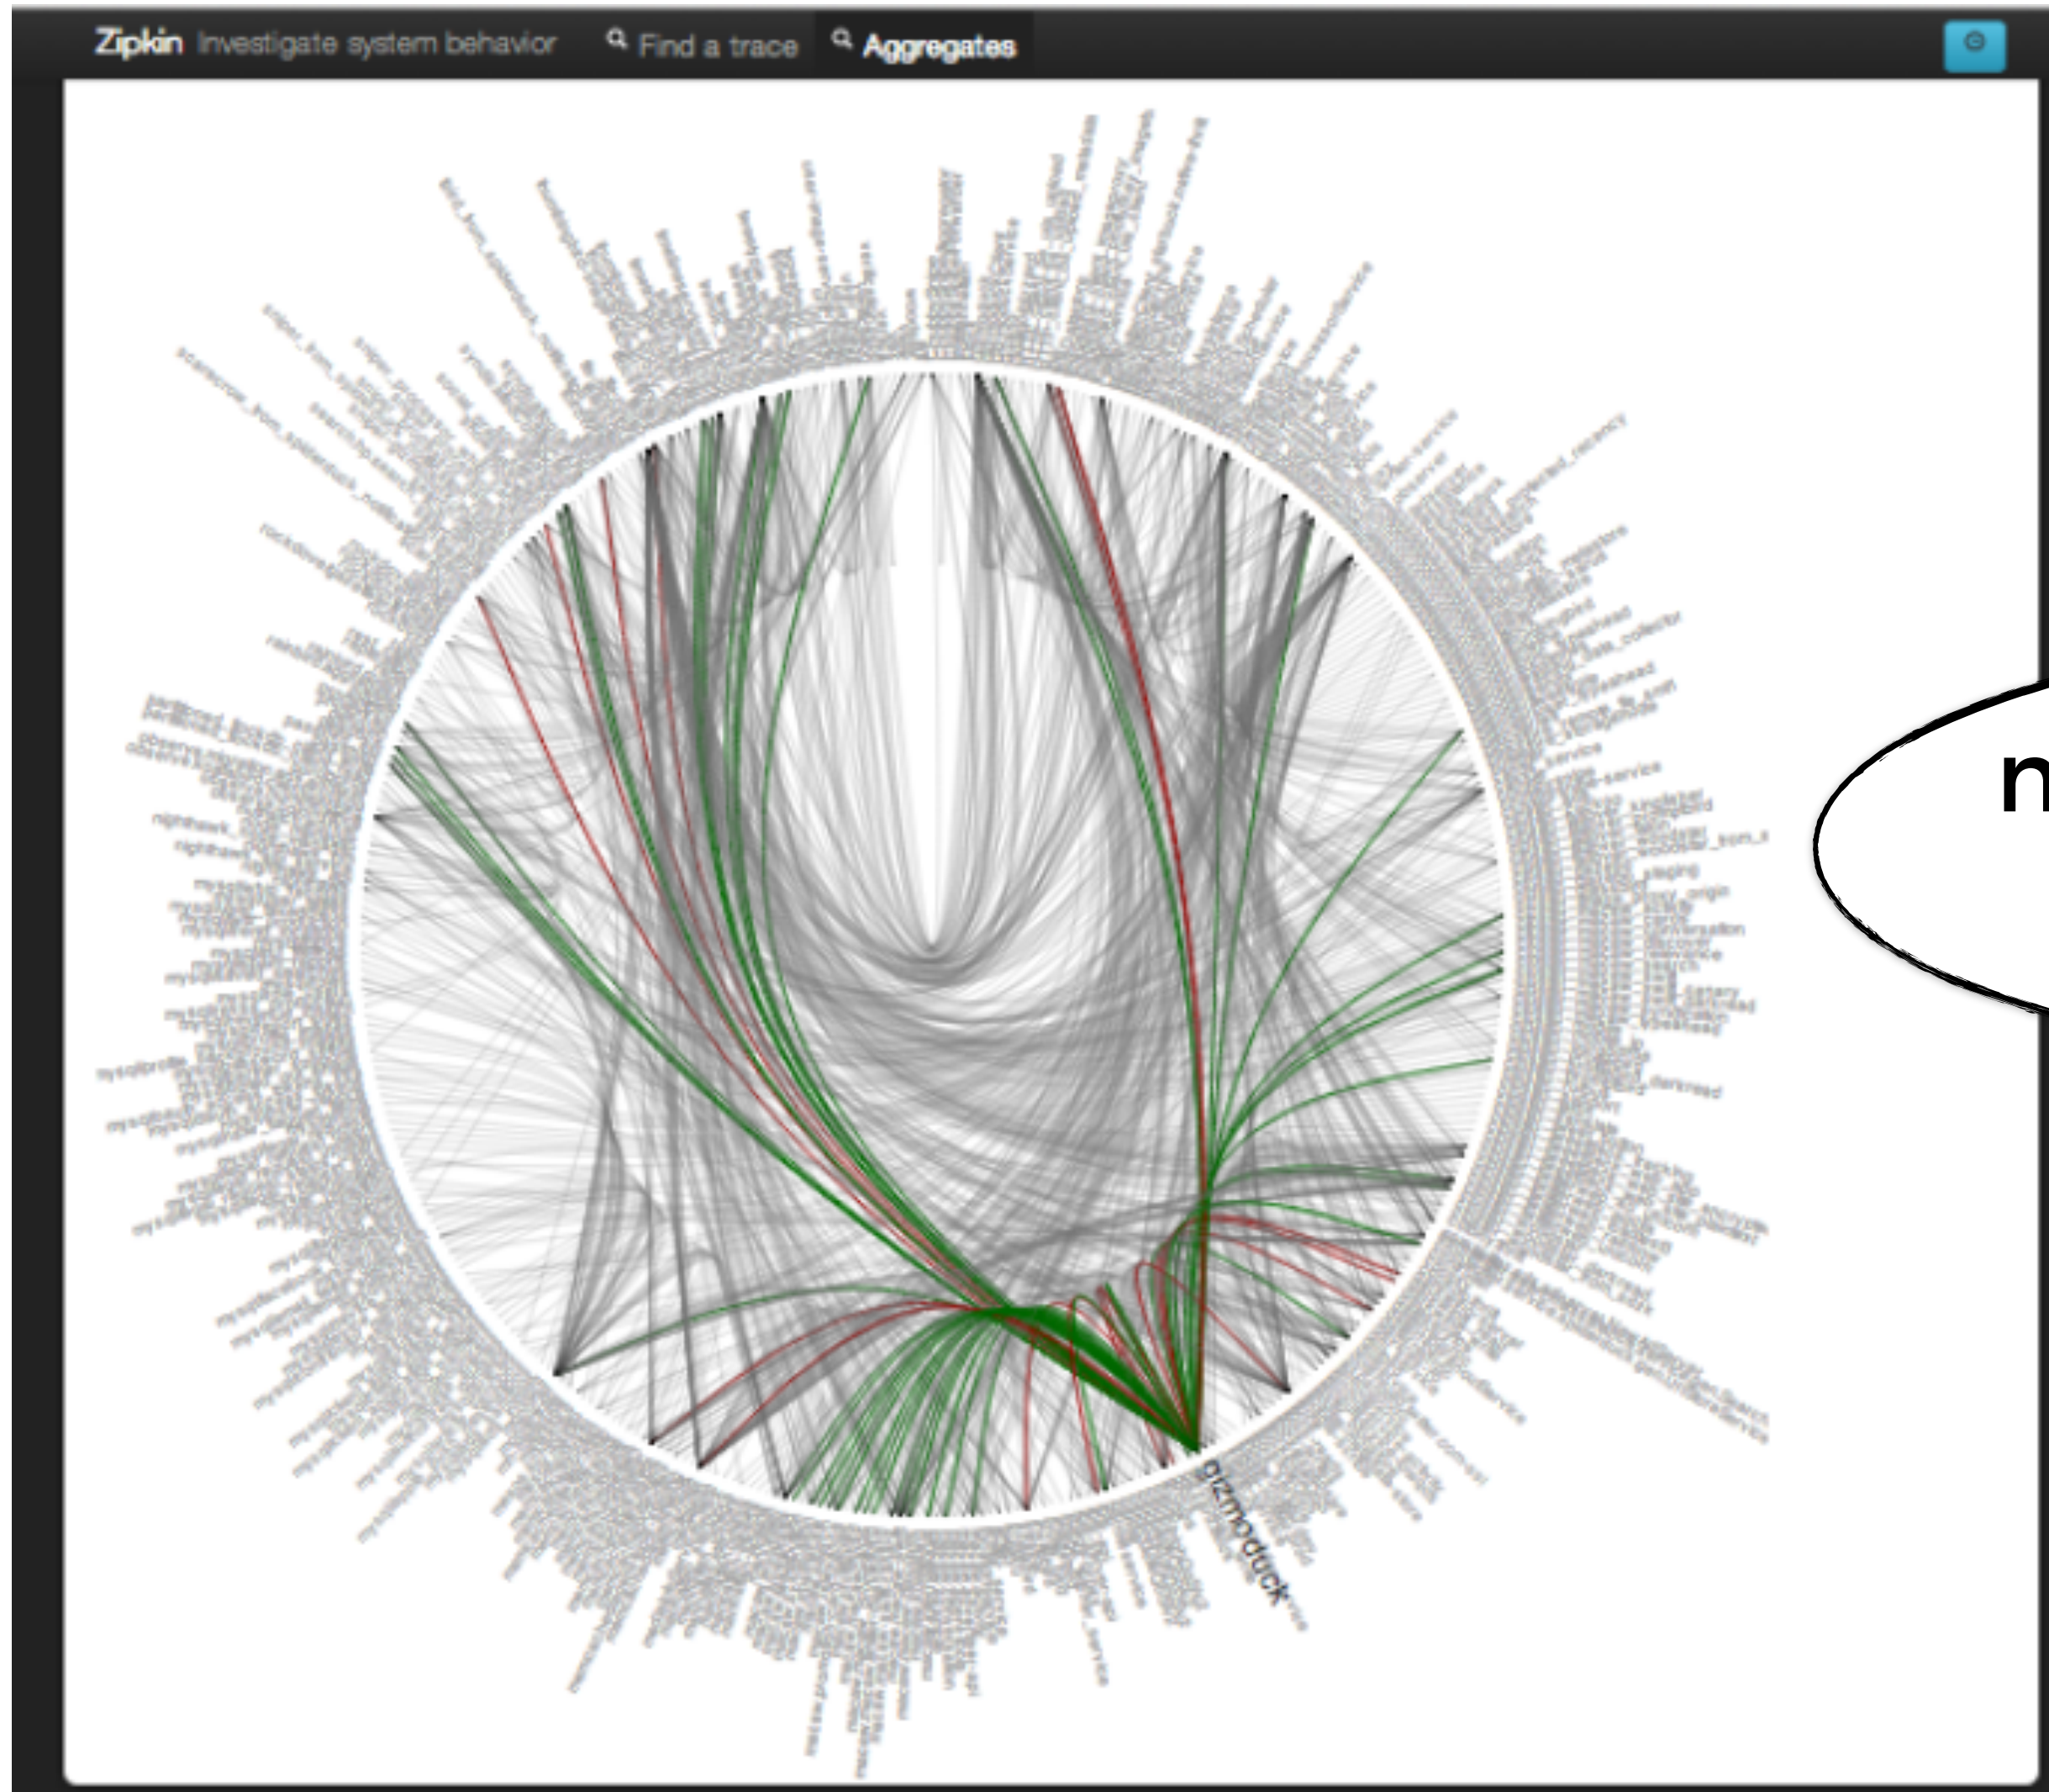
f in ↗

As Twitter has moved from a monolithic to a distributed architecture, our scalability has increased dramatically.

Because of this, the overall complexity of systems and their interactions has also escalated. This decomposition has led to Twitter managing hundreds of services across our datacenters. Visibility into the health and performance of our diverse service topology has become an important driver for quickly determining the root cause of issues, as well as increasing Twitter's overall reliability and efficiency. Debugging a complex program might involve instrumenting certain code paths or running special utilities; similarly Twitter needs a way to perform this sort of debugging for its distributed systems.



<http://bit.ly/2EXKEFZ>



mm... I wonder what's going on here...



what is **observability**?
how is it different from monitoring?

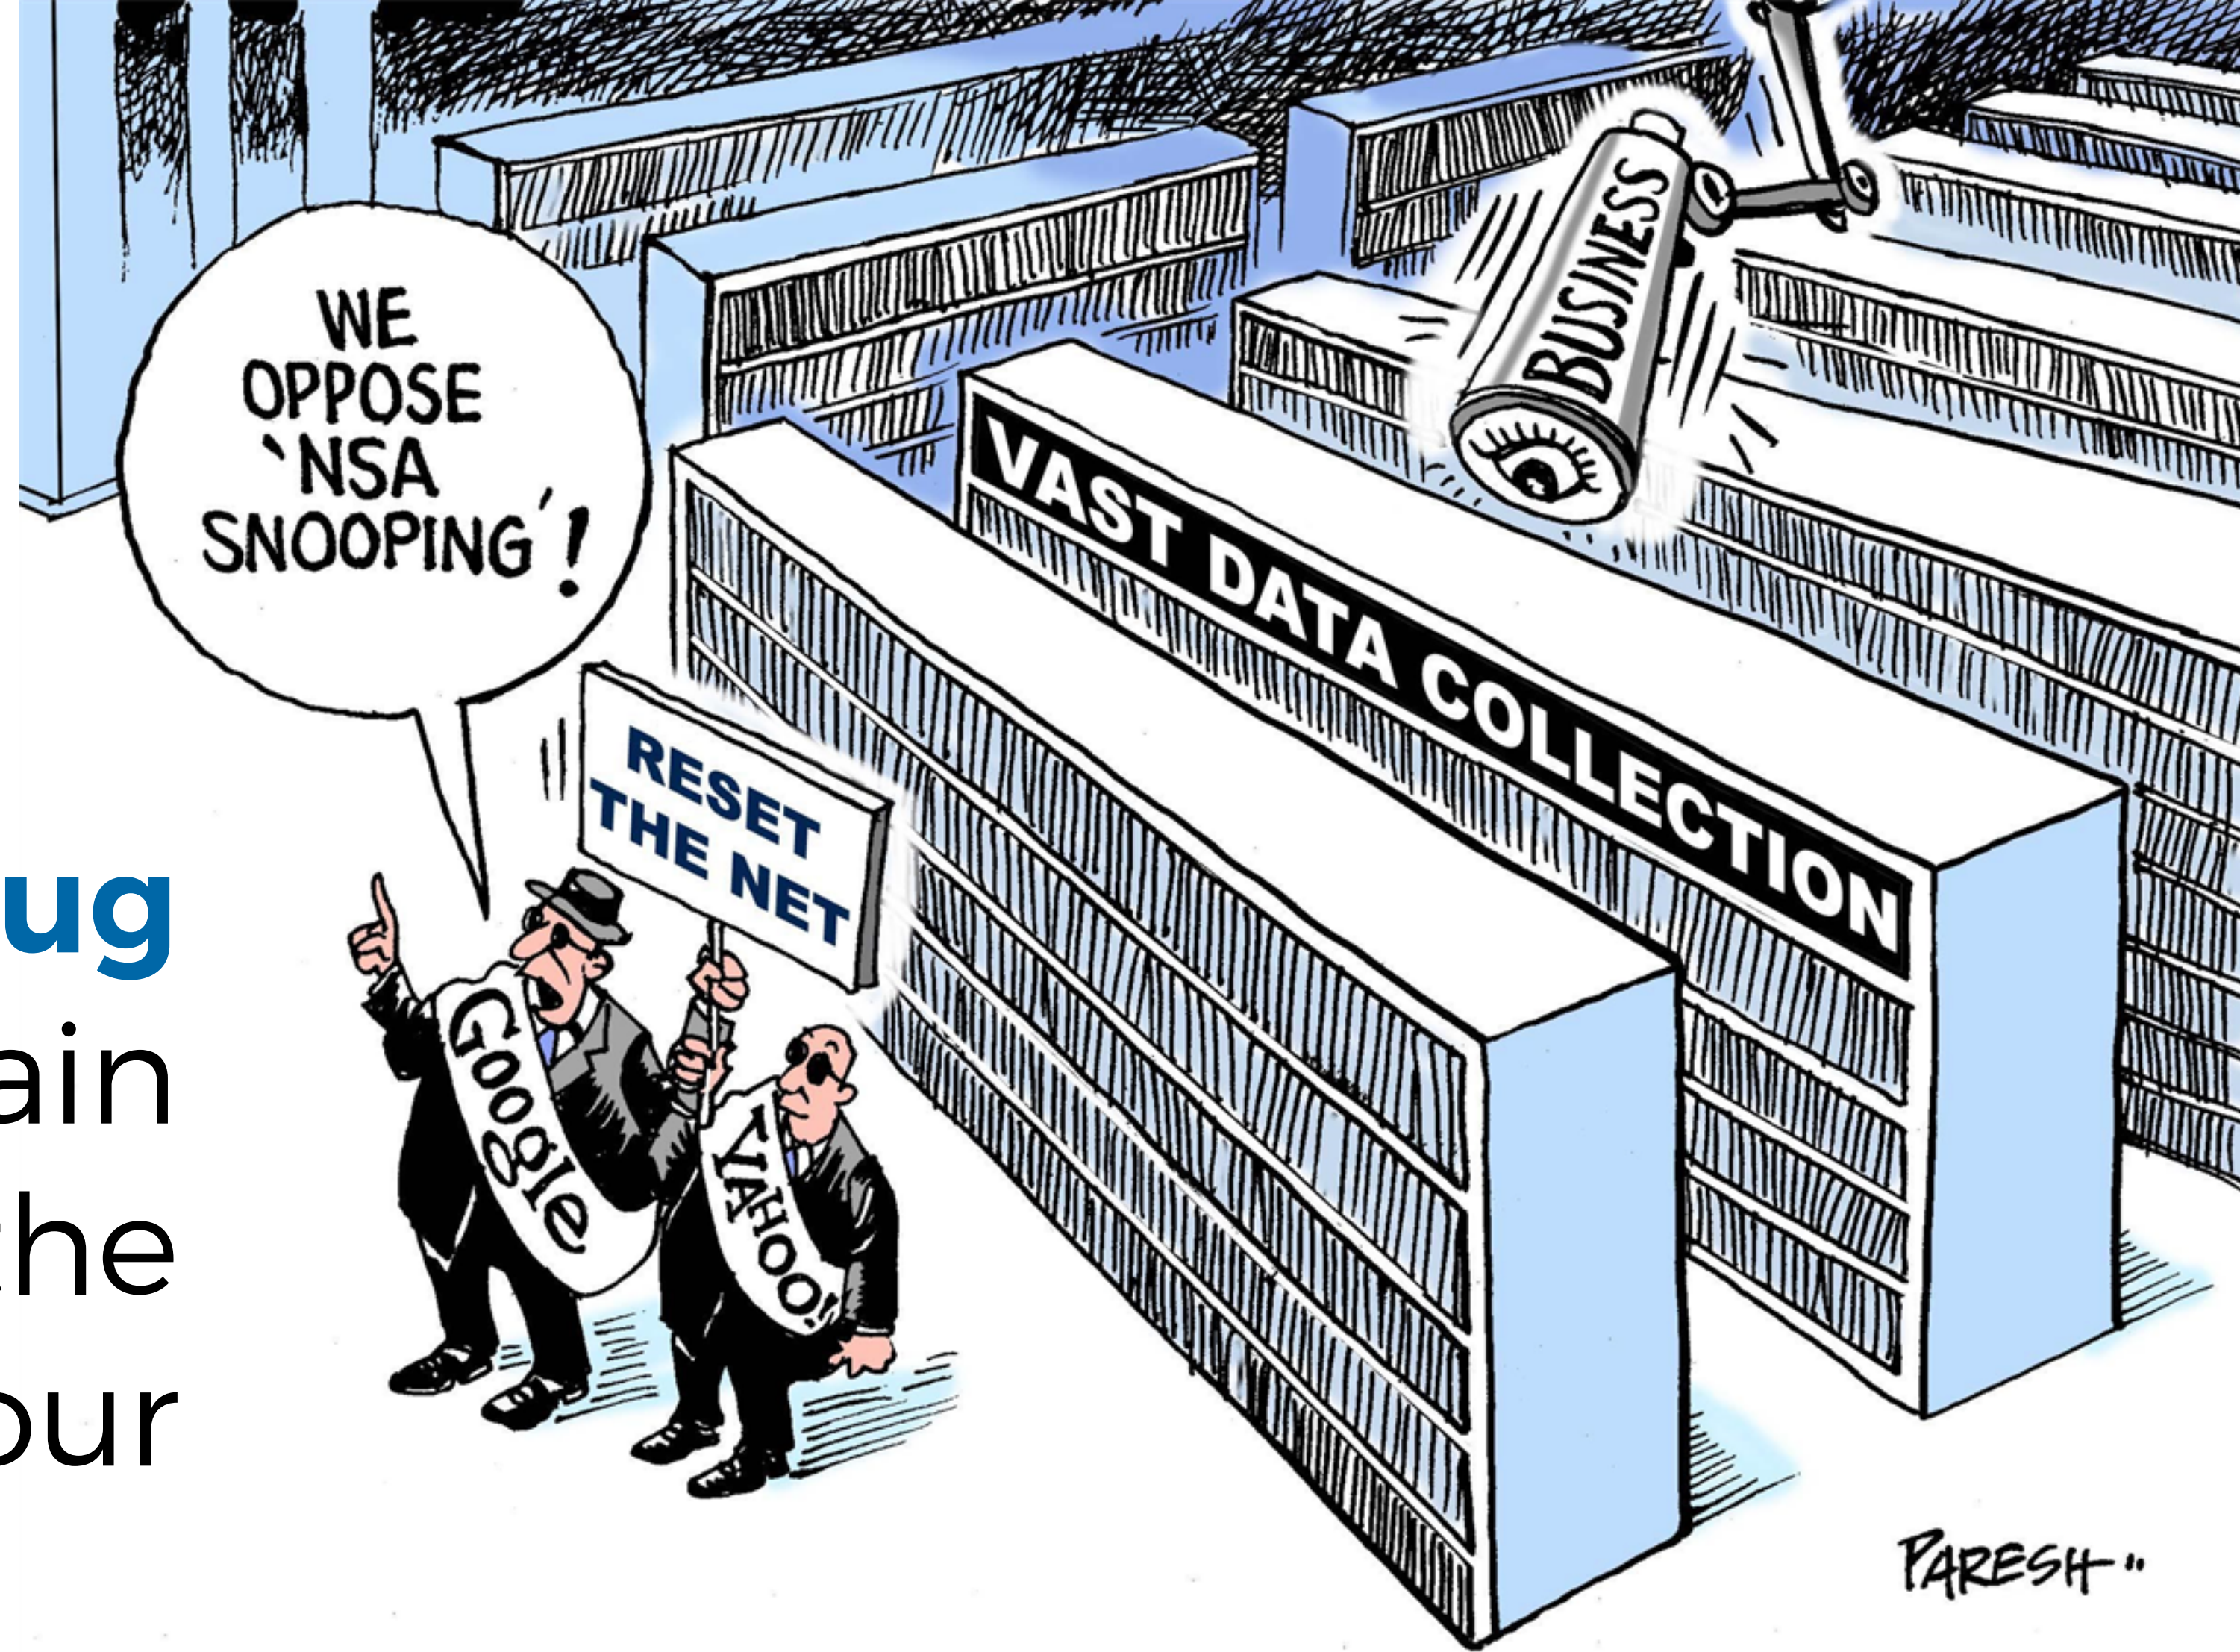




Monitoring

watching out for **known** failure modes in the system, e.g. network I/O, CPU, memory usage, ...

being able to **debug** the system, and gain **insights** into the system's behaviour



PARESH"

Cagle.com

Observability

“

*However, I would argue that the health of the system no longer matters. We've entered an era where what matters is the health of each individual event, or each individual user's experience, or each shopping cart's experience (or other high cardinality dimensions). **With distributed systems you don't care about the health of the system, you care about the health of the event or the slice.** ”*

- Charity Majors <http://bit.ly/2E2QngU>


“

*However, I would argue that the health of the system no longer matters. We've entered an era where what matters is the health of each individual event, or each individual user's experience, or each shopping cart's experience (or other high cardinality dimensions). **With distributed systems you don't care about the health of the system, you care about the health of the event or the slice.** ”*

- Charity Majors <http://bit.ly/2E2QngU>



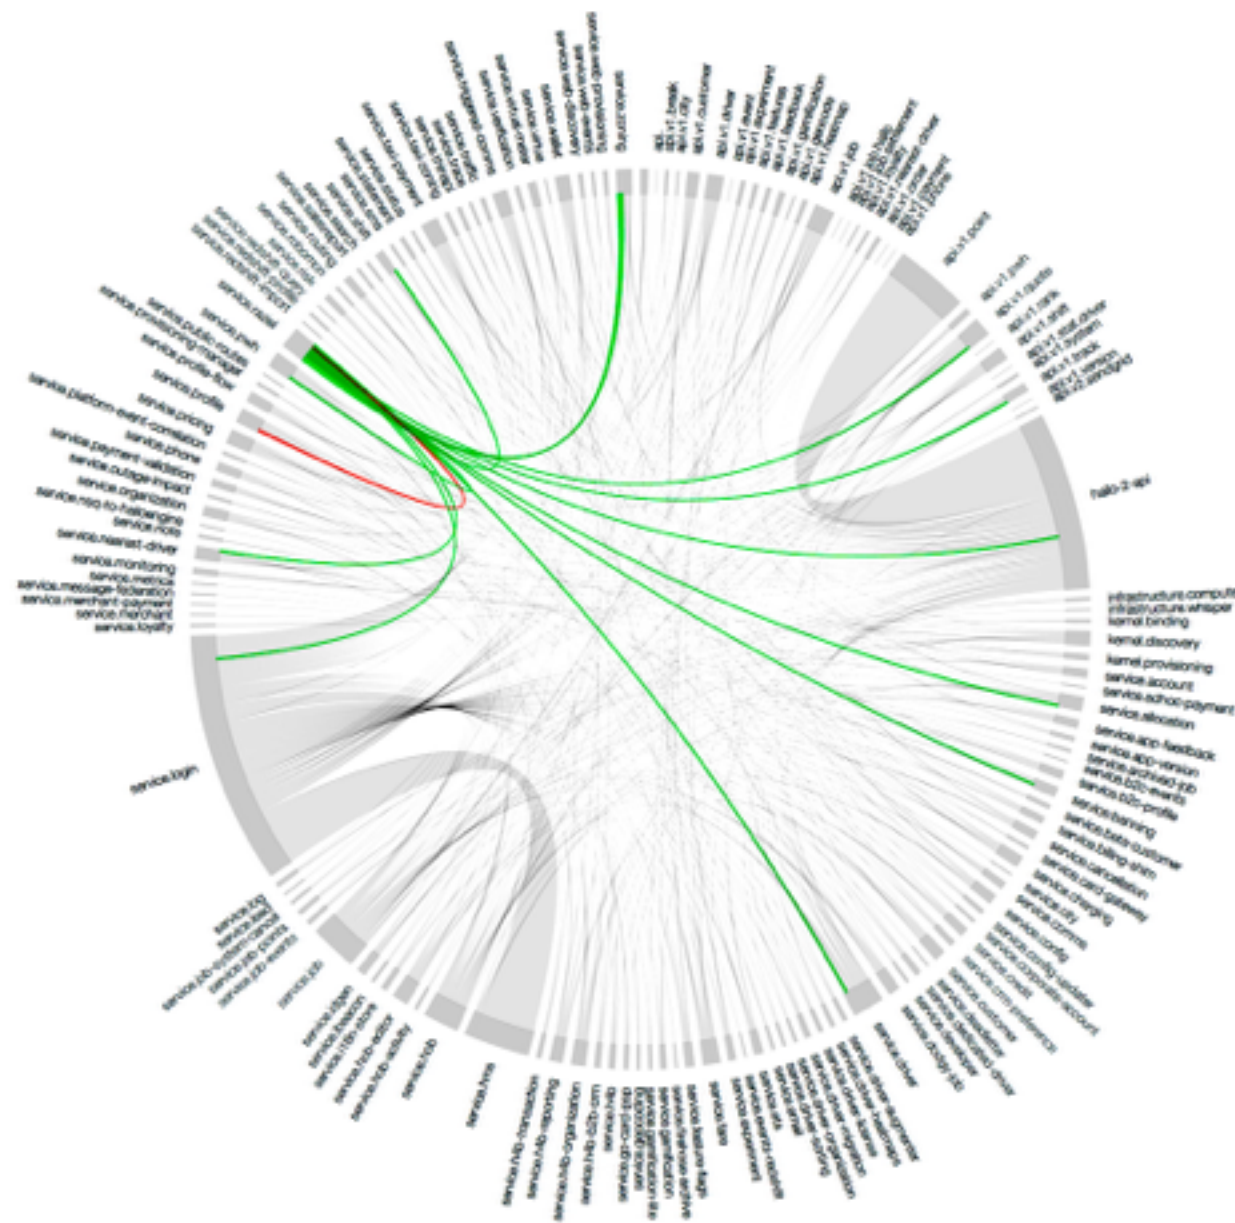
These are the four pillars of the Observability Engineering team's charter:

- *Monitoring*
- *Alerting/visualization*
- *Distributed systems tracing infrastructure*
- *Log aggregation/analytics* 

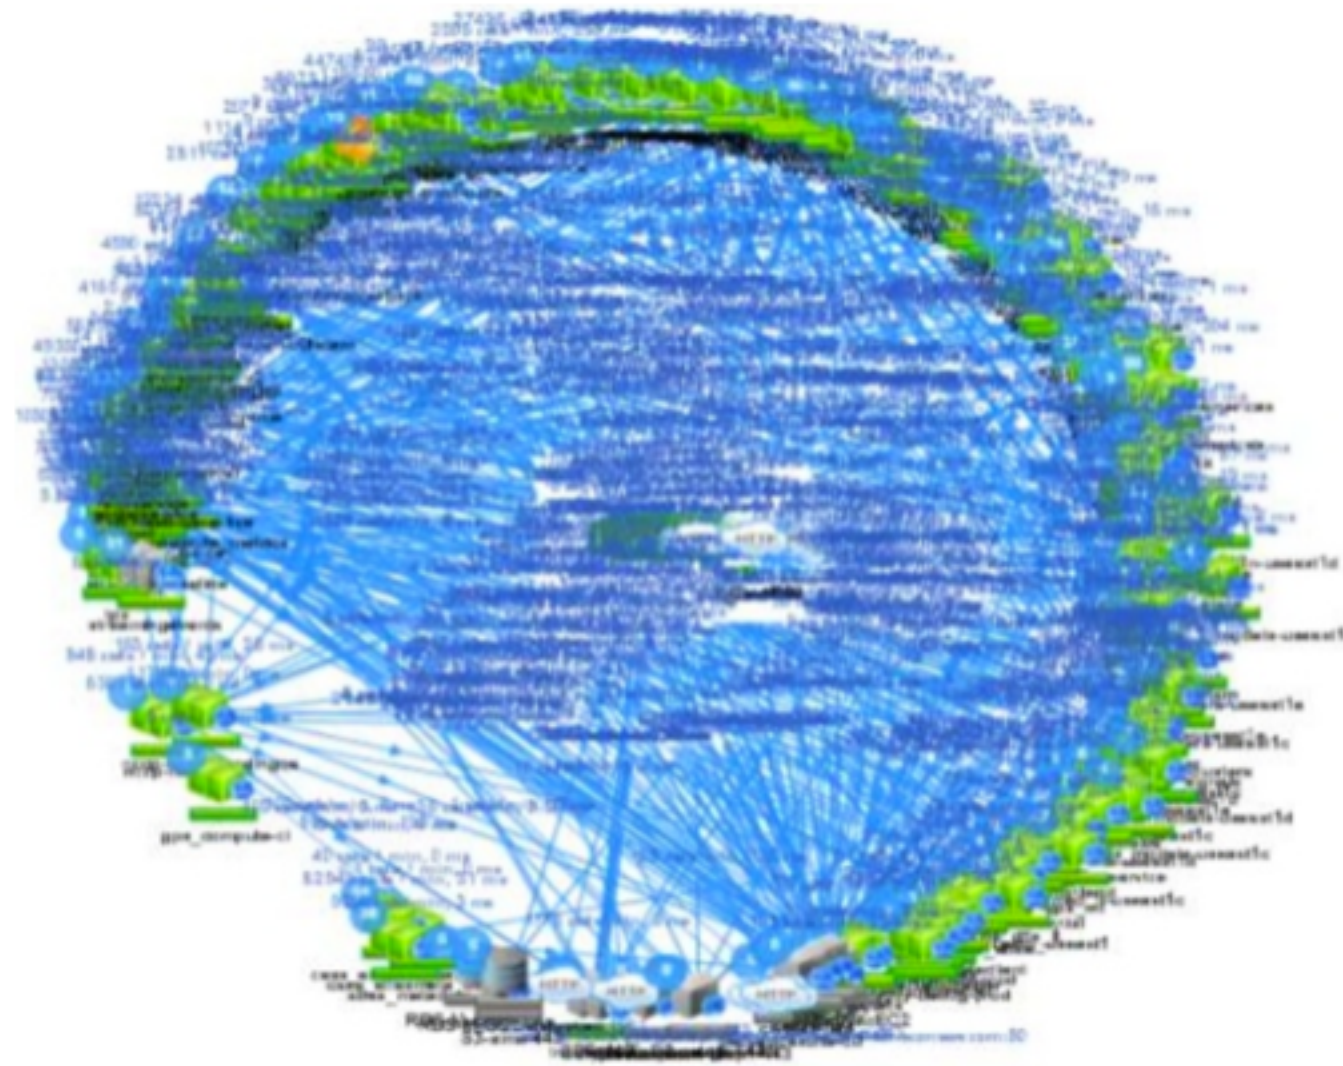
- Observability Engineering at Twitter <http://bit.ly/2DnjyuW>

Observability is useful even outside of incidents and outages

450+ microservices

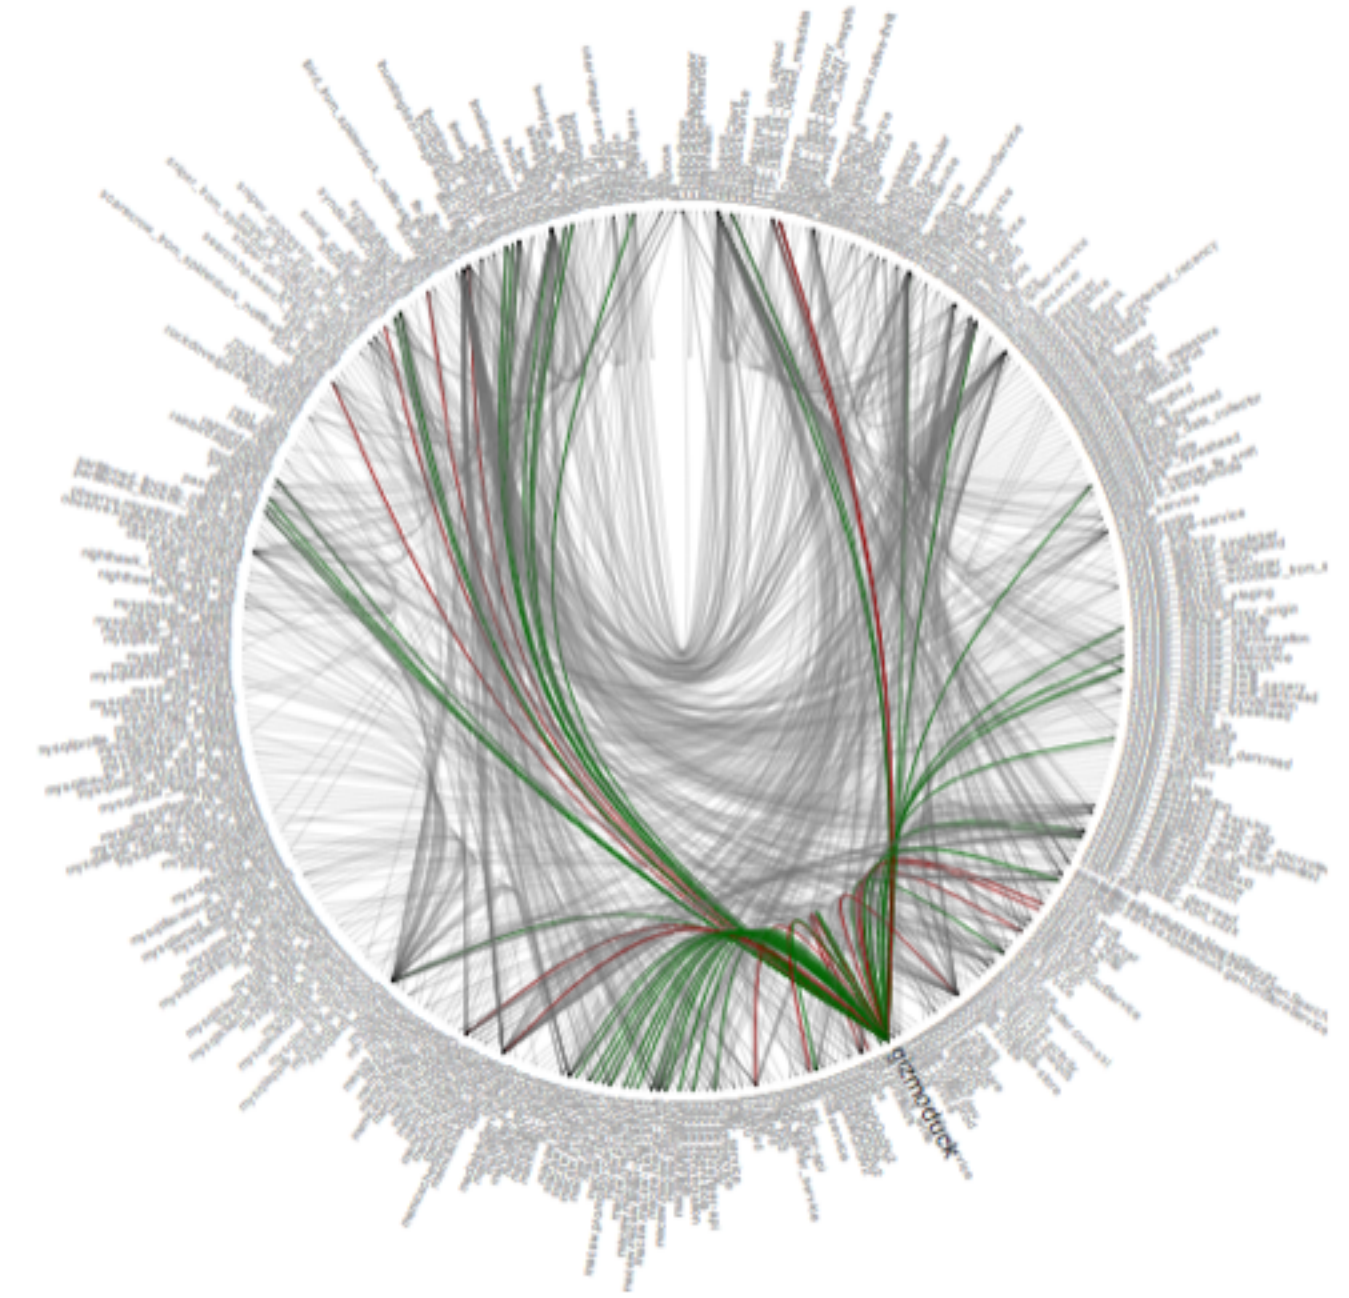


500+ microservices



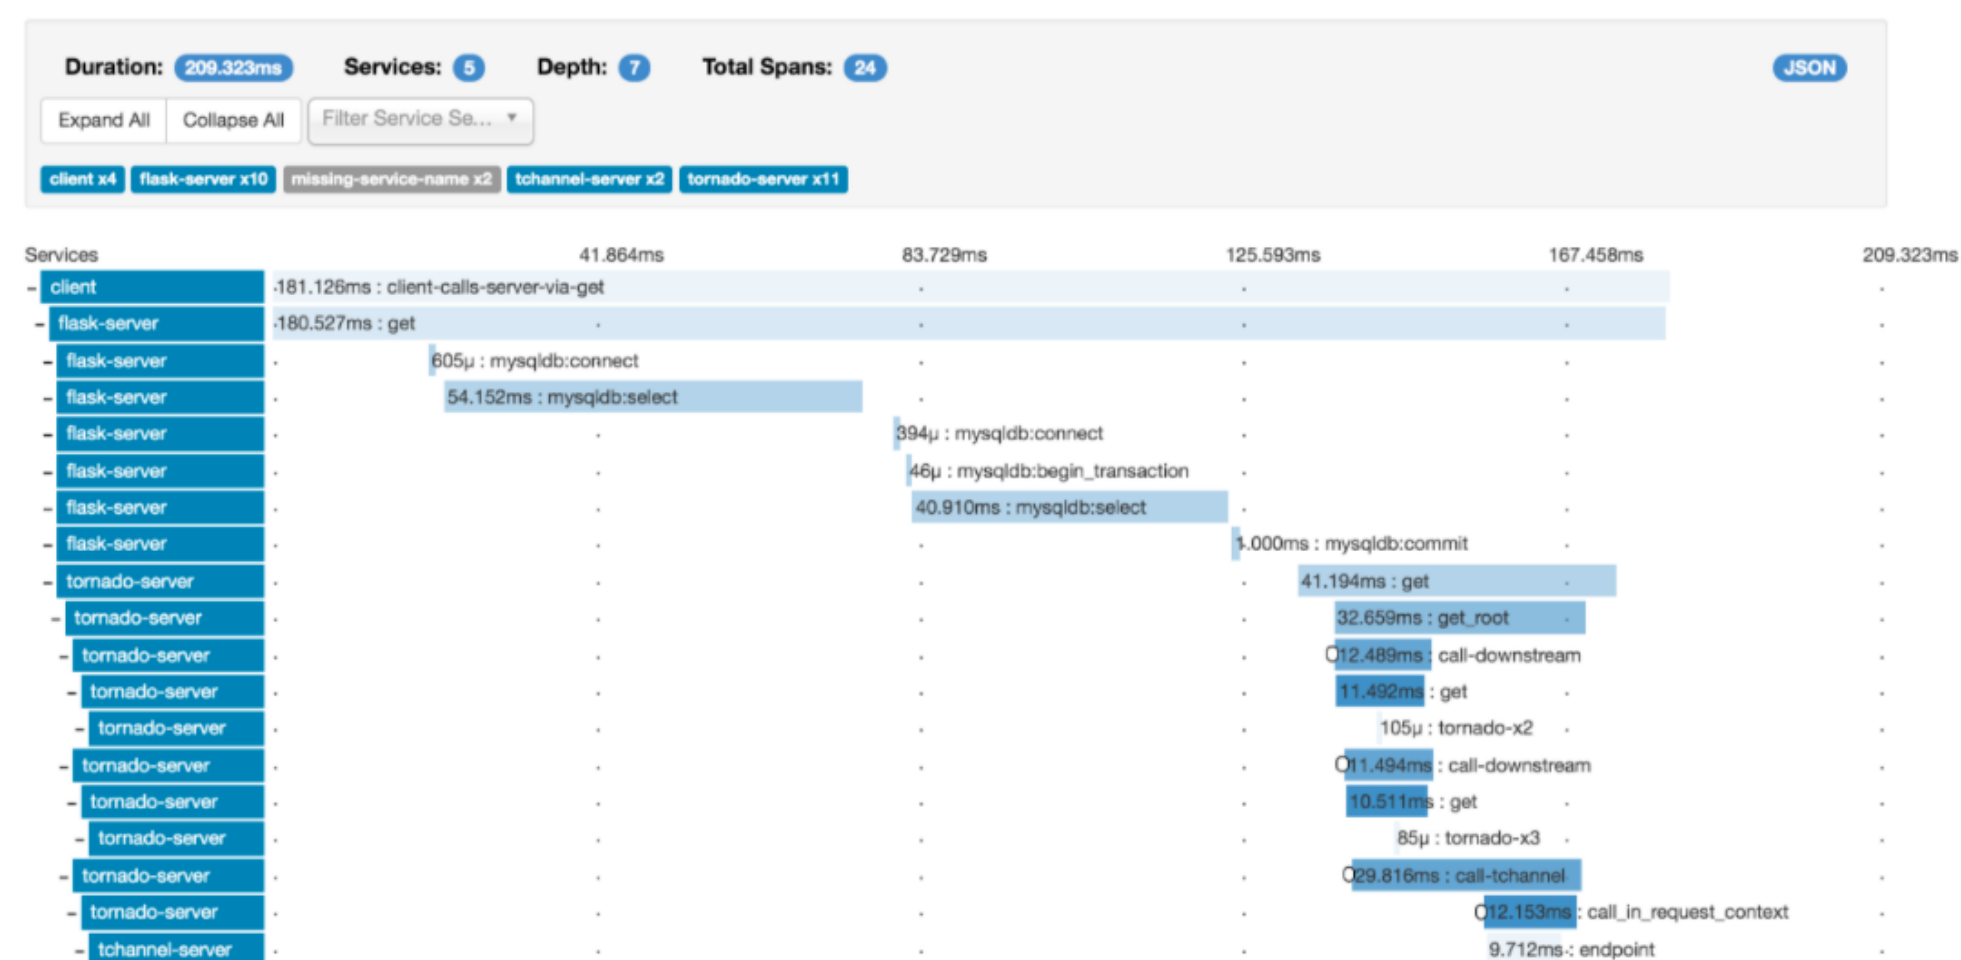
NETFLIX

500+ microservices



microservices death stars circa 2015

Zipkin



Zipkin is a distributed tracing system. It helps gather timing data needed to troubleshoot latency problems in microservice architectures. It manages both the collection and lookup of this data. Zipkin's design is based on the [Google Dapper](#) paper.

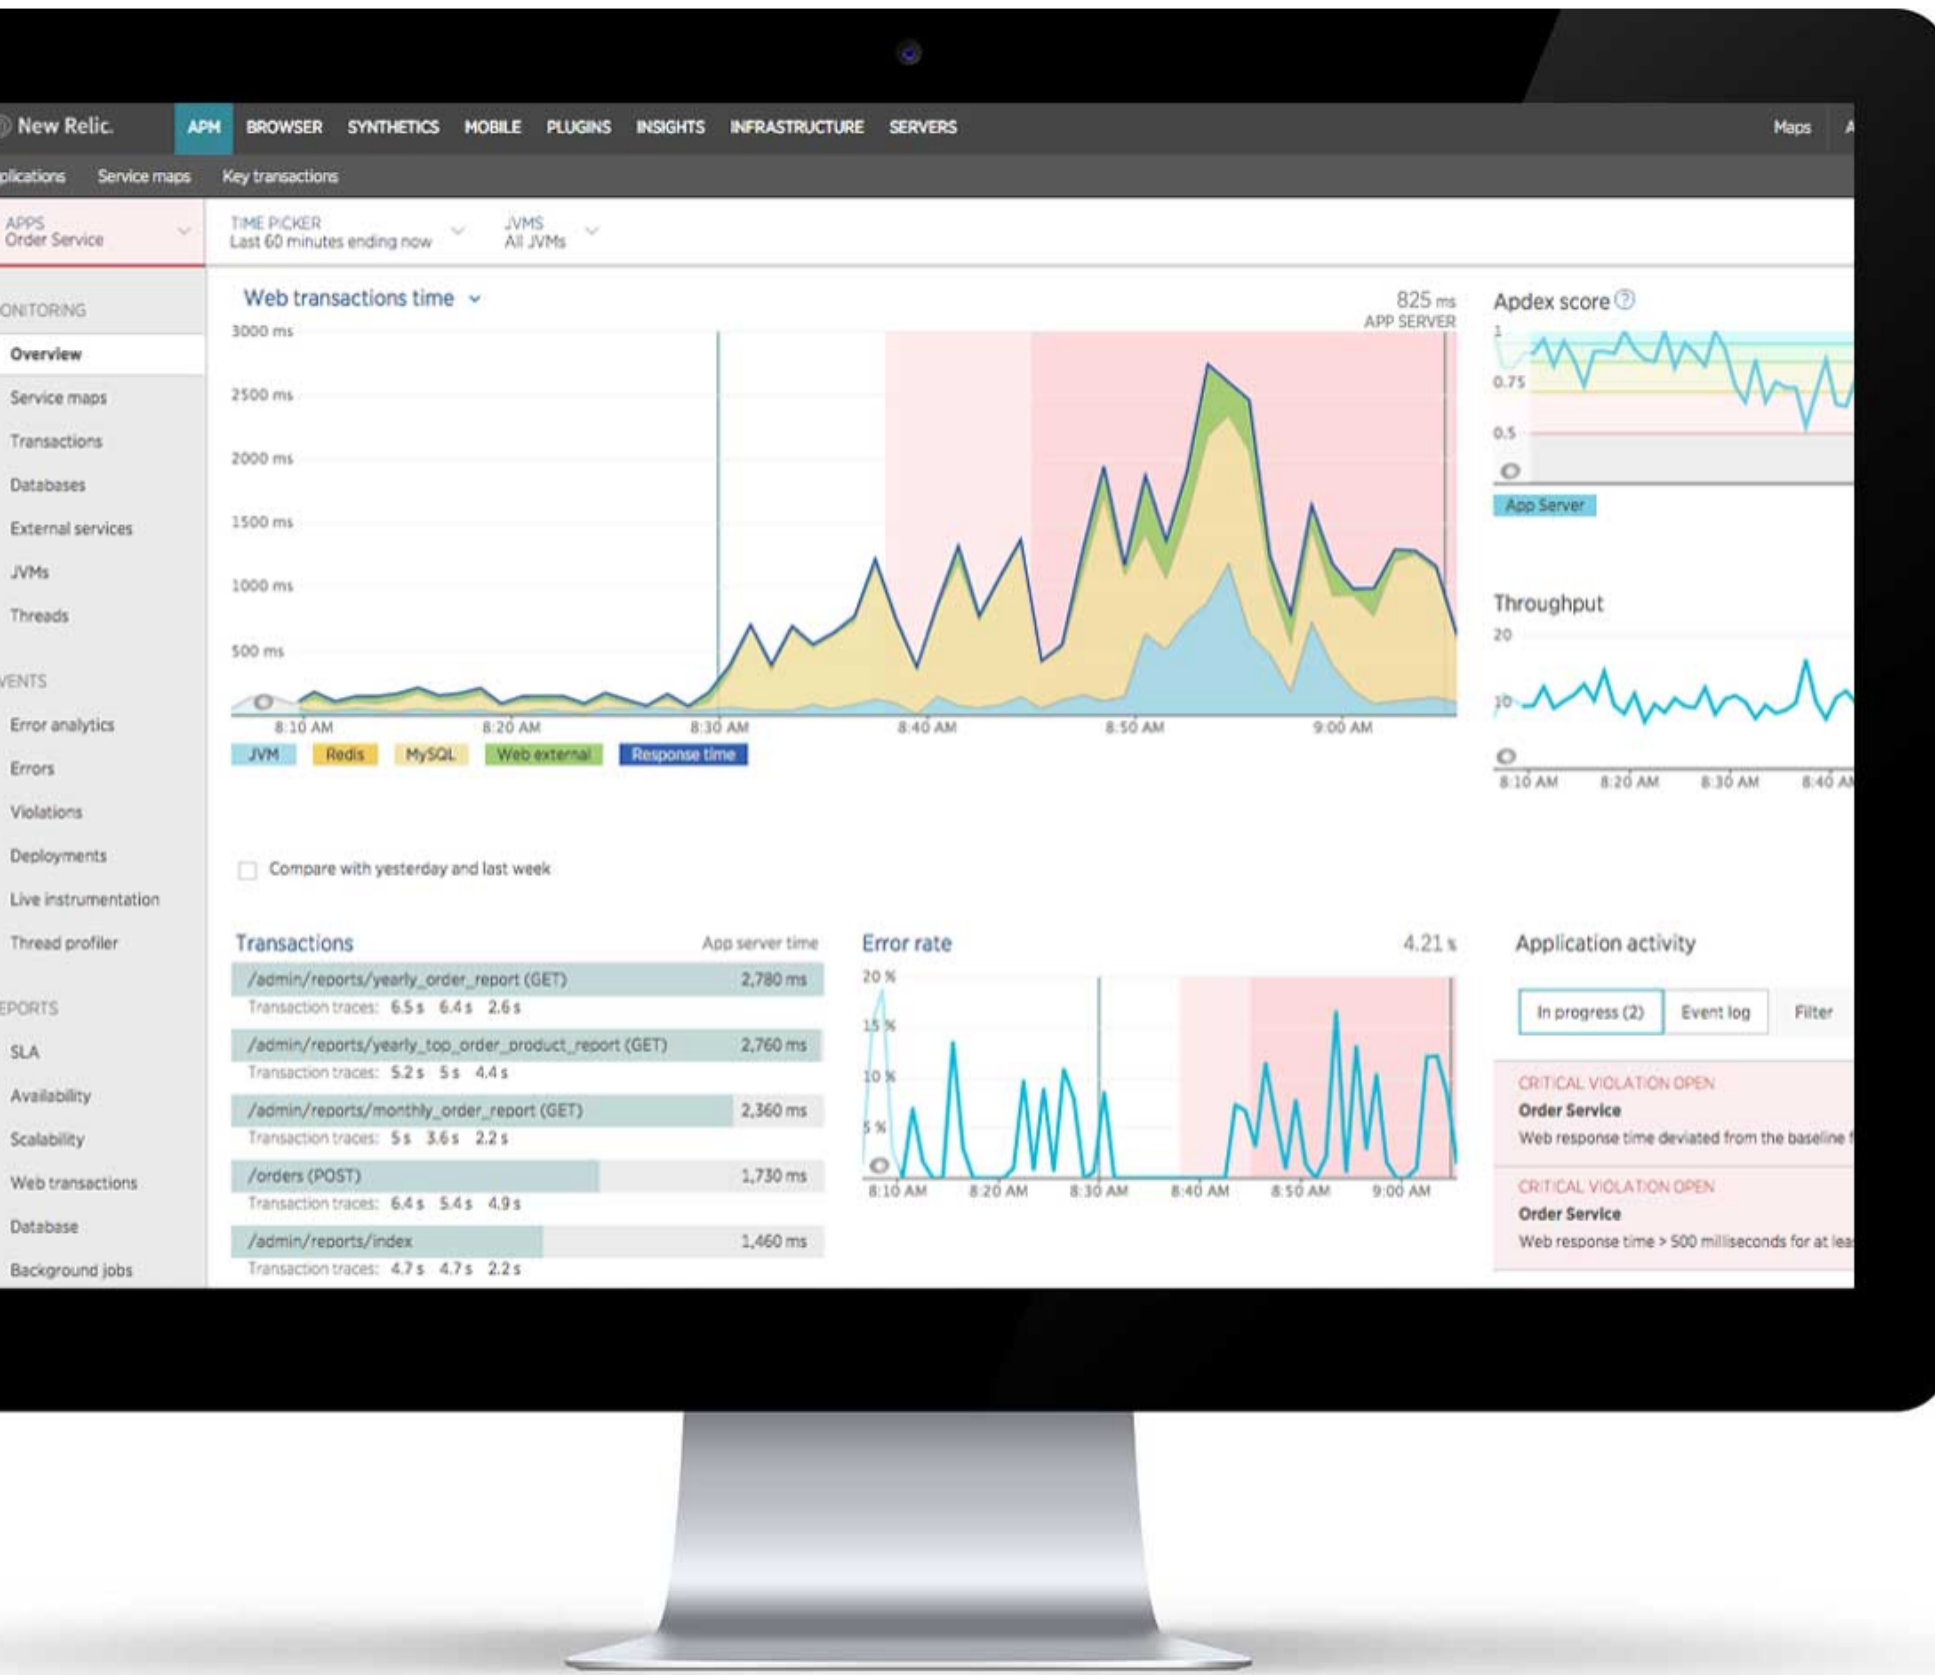
Applications are instrumented to report timing data to Zipkin. The Zipkin UI also presents a Dependency diagram showing how many traced requests went through each application. If you are troubleshooting latency problems or errors, you can filter or sort all traces based on the application, length of trace, annotation, or timestamp. Once you select a trace, you can see the percentage of the total trace time each span takes which allows you to identify the problem application.

Where to go next?

- To try out Zipkin, check out our [Quickstart guide](#)
- See if your platform has an [existing instrumentation library](#)
- Join the [Zipkin Gitter chat channel](#)
- The source code is on GitHub as [openzipkin/zipkin](#)
- Issues are also tracked on [GitHub](#)



- Home
- Quickstart
- Architecture
- Existing instrumentations
- Zipkin Community
- Data Model
- Instrumenting a library



Application performance magic.

Build, deploy, and maintain great software with New Relic APM.

Request a demo

NEW Distributed Tracing

Distributed Tracing support will let you track a user request through

It's easy to stay up-to-date.

Your Email Address: *

Input field for email address





Observability for a distributed world

YOUR CORP ▾ **MULTISERVICE WEBS** ▾ **QUERIES** OVERVIEW SAMPLES ▾ DOCS

BREAK DOWN *None; use all rows* | **CALCULATE** AVG(response_time_ms) MAX(response_time_ms) | **FILTER** hostname = app1 | **ORDER** AVG(response_time_ms) desc | **LIMIT** None | **Run Query** Run a few seconds ago

Query at 9/6 7:06PM [\(Click to edit\)](#)

Last 2 hr → ABS | [★ Add to Playlist](#) | [👤 Share](#) | [📄 View Raw Data](#) | [⚙️ Graph Settings](#)

AVG(response_time_ms)

Deploy #047

MAX(response_time_ms)

My History | **Team History**

Once you've run some queries, your own results will appear here. In the meantime, here's what your teammates have been up to:

TODAY ▾

Run at 7:06P

- T 5:06P - 7:06P
- C AVG(response_time_ms) MAX(response_time_ms)
- F hostname = app1
- O AVG(response_time_ms) de

Run at 7:06P

- T 5:06P - 7:06P
- C AVG(response_time_ms) MAX(response_time_ms)
- B hostname
- O AVG(response_time_ms) de

Run at 7:06P

- T 5:06P - 7:06P
- C AVG(response_time_ms)
- B hostname
- O AVG(response_time_ms) de





TransAlta Teams With Splunk for Security and Operational Intelligence

Saves up to \$1 million while reducing investigation times from days to minutes

[Read the Case Study](#)

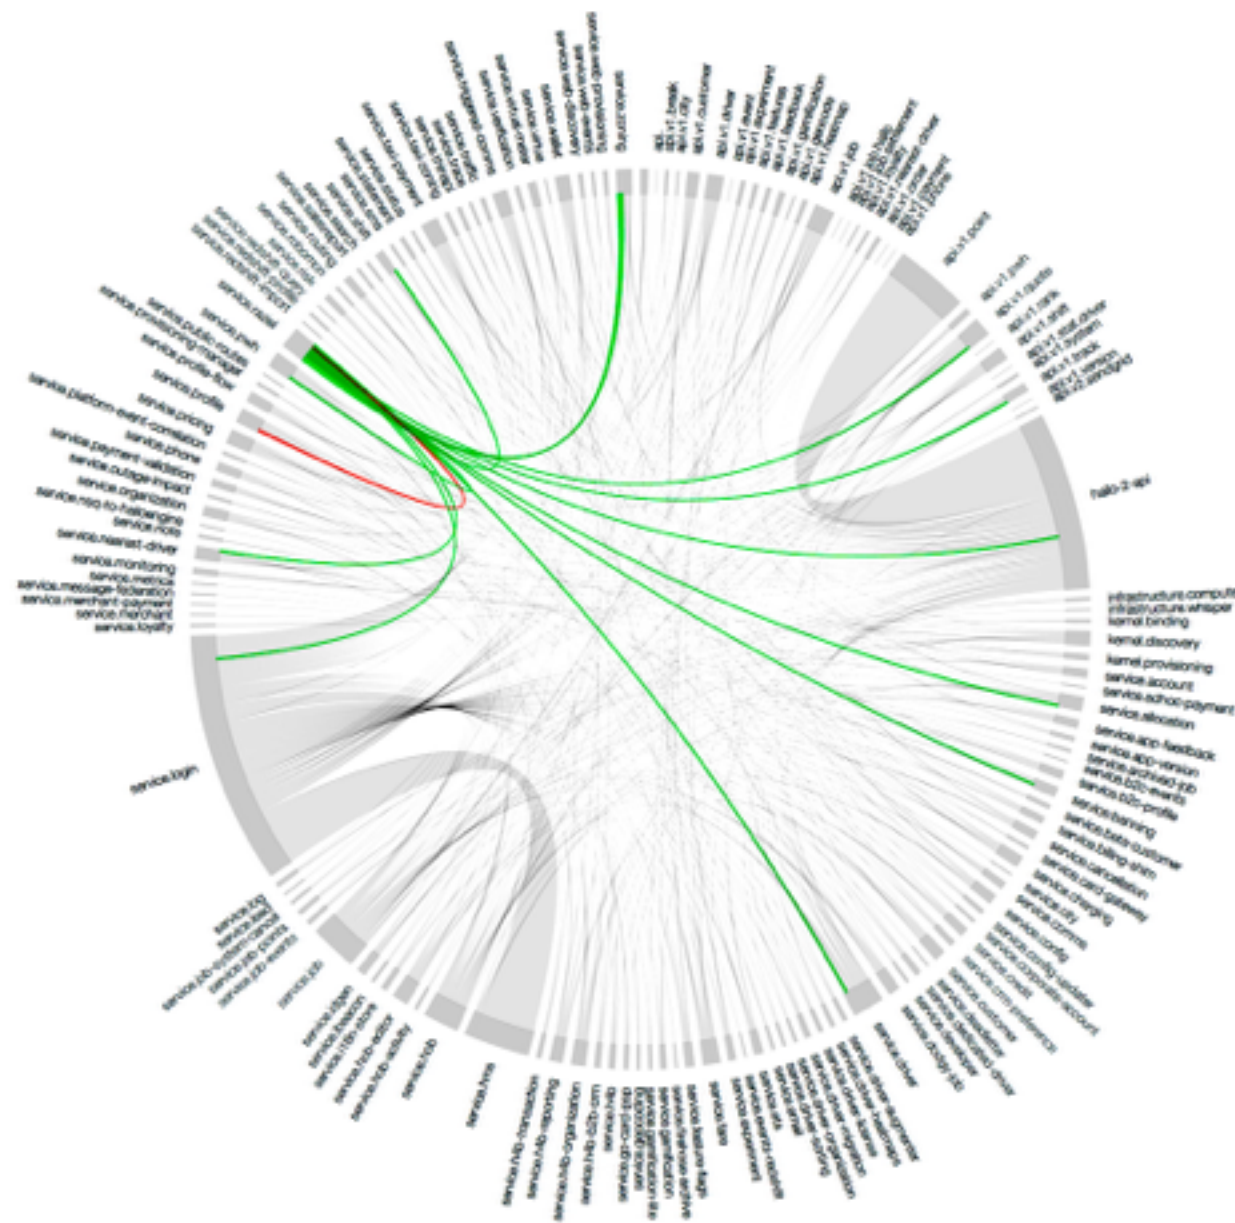
Gartner names Splunk a SIEM Magic Quadrant leader for the fifth year running!

[Read the Report](#)

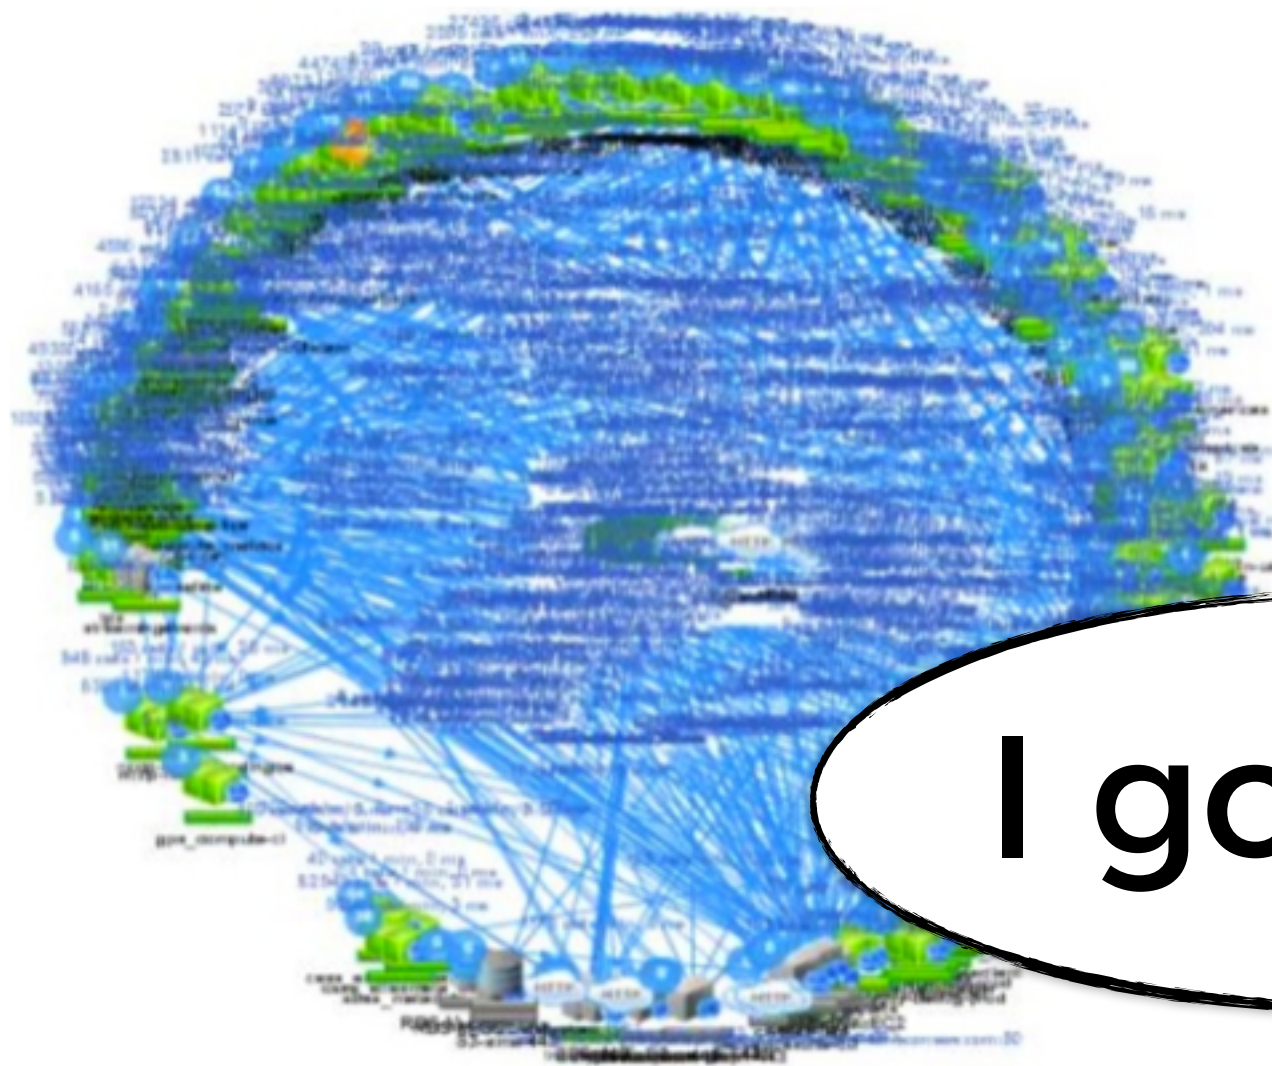
Turn Machine Data Into Answers



450+ microservices

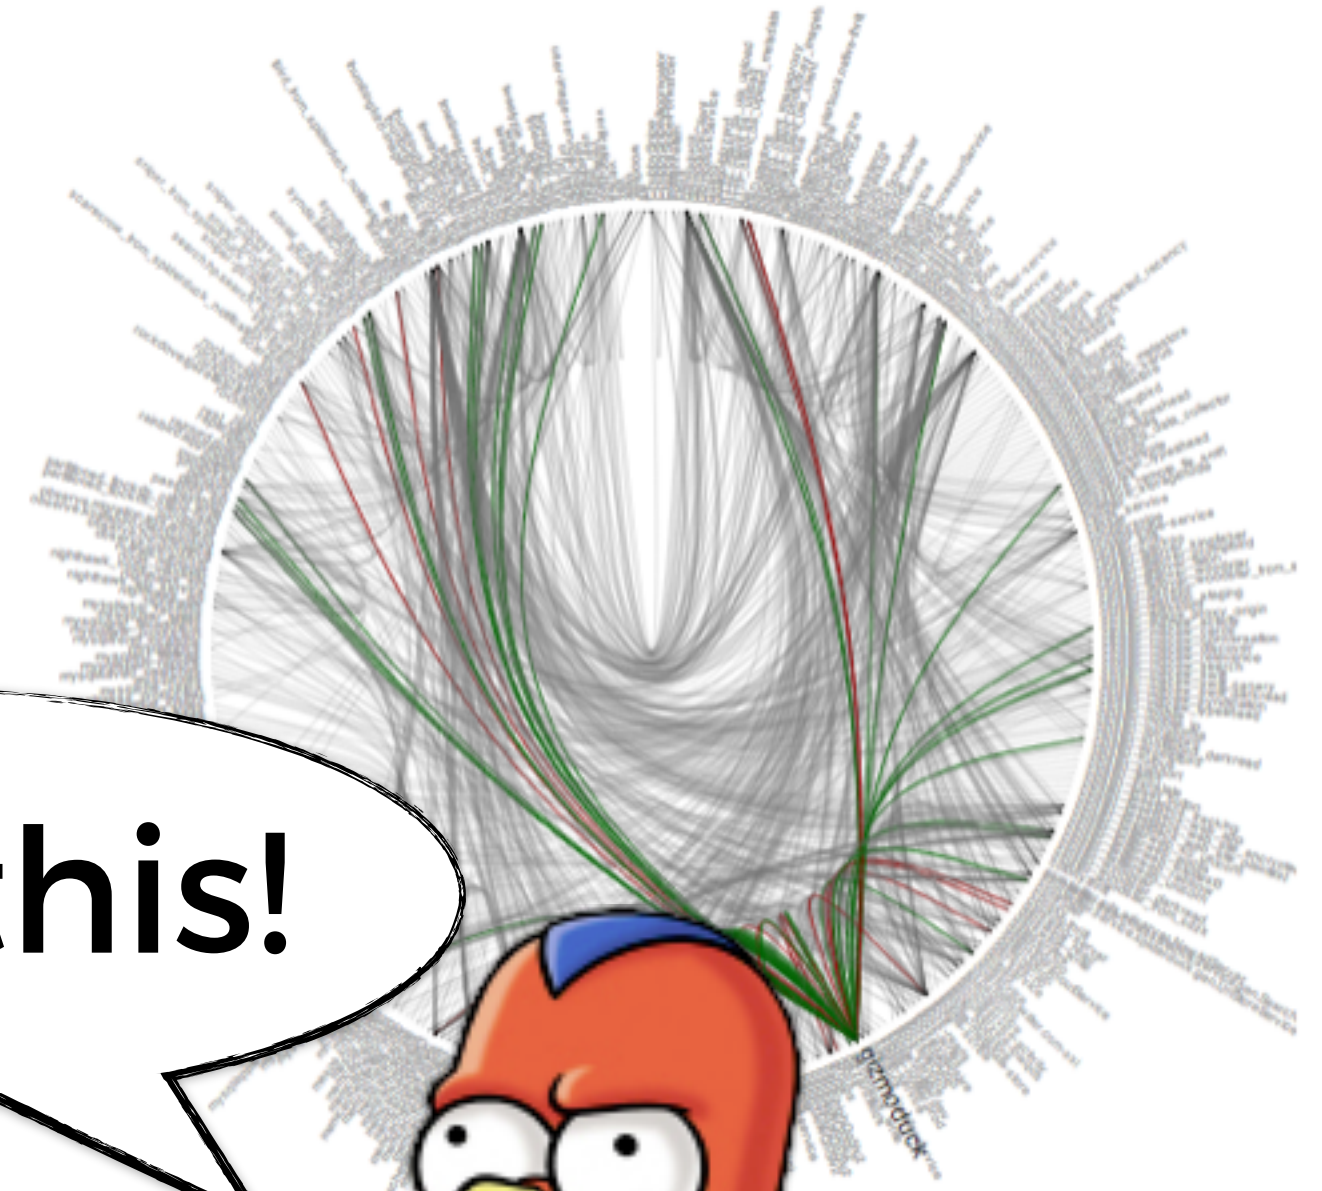


500+ microservices



NETFLIX

500+ microservices



I got this!



microservices death stars circa 2011

Introducing AWS Lambda

An event-driven computing service for
dynamic applications



new

challenges



new

challenges



NO ACCESS
to underlying OS





NOWHERE

to install agents/daemons

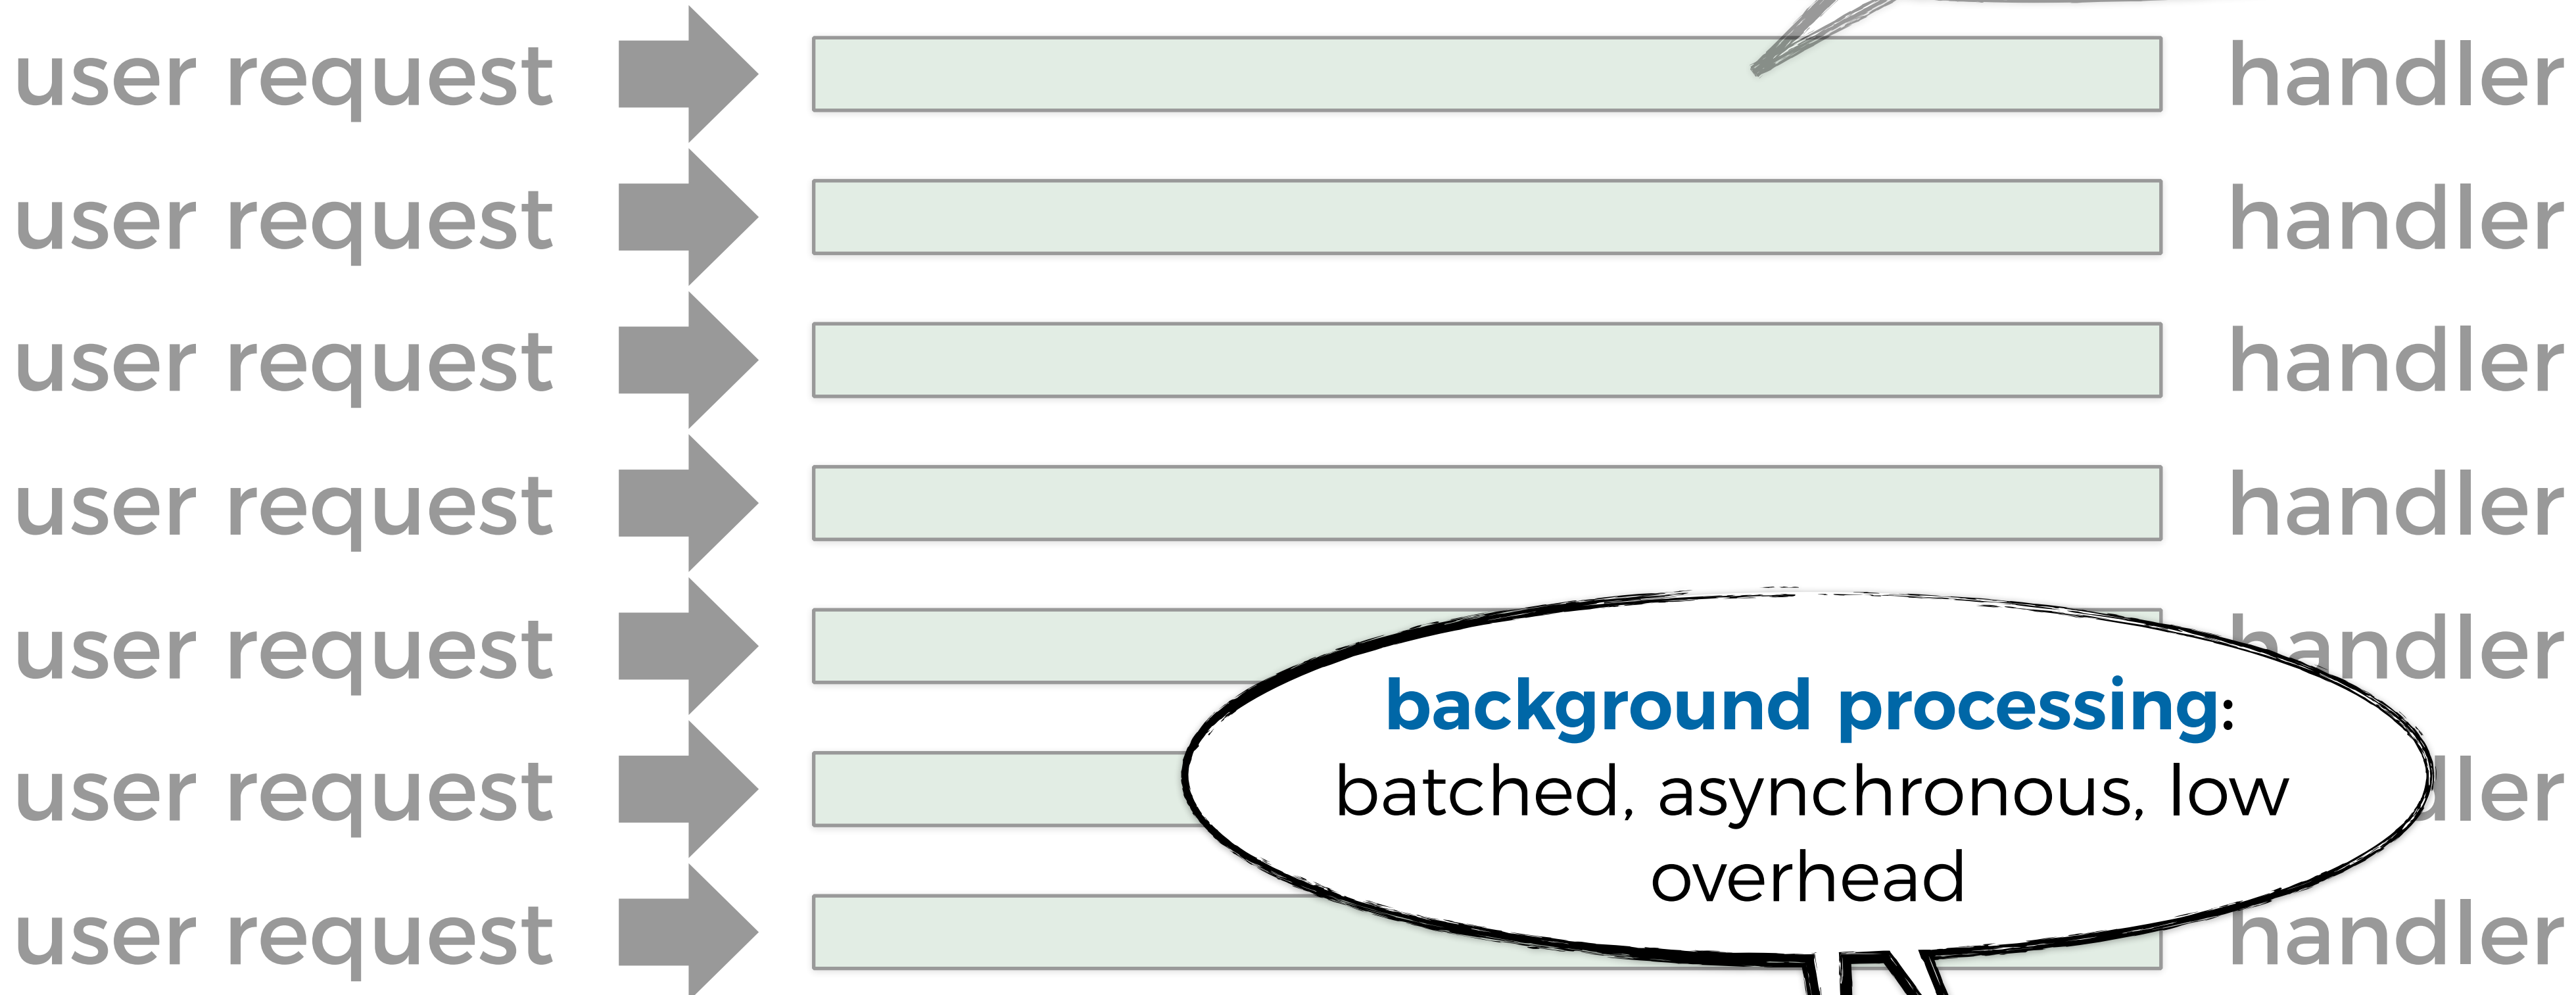
new challenges

- **nowhere** to install agents/daemons

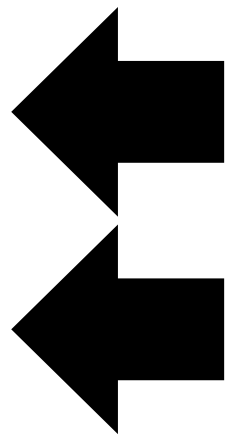
critical paths:
minimise user-facing latency



critical paths:
minimise user-facing latency



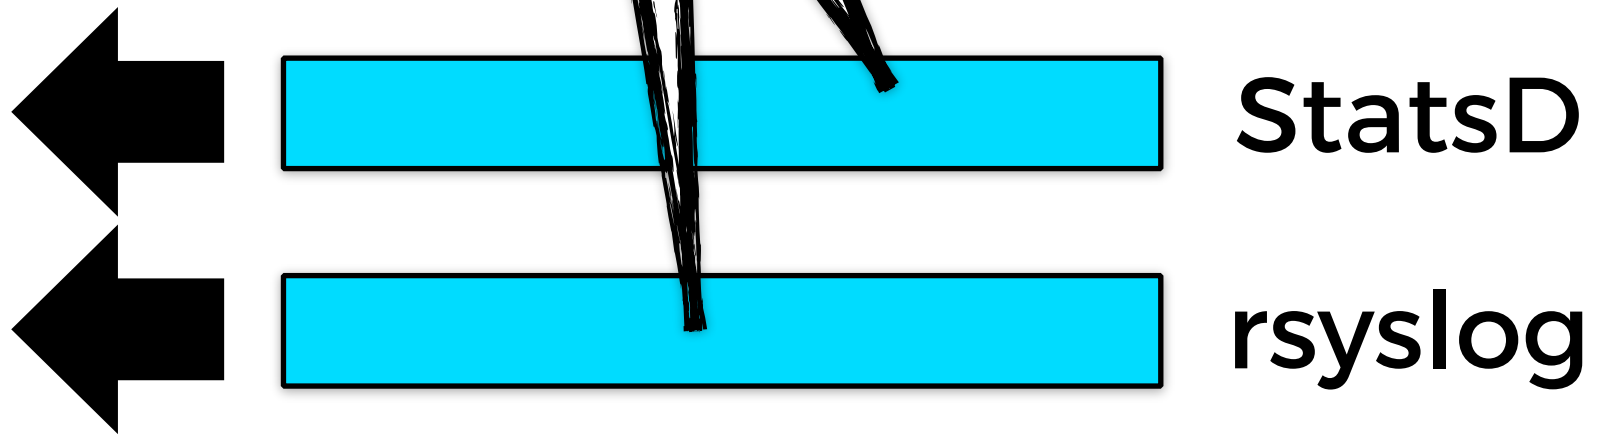
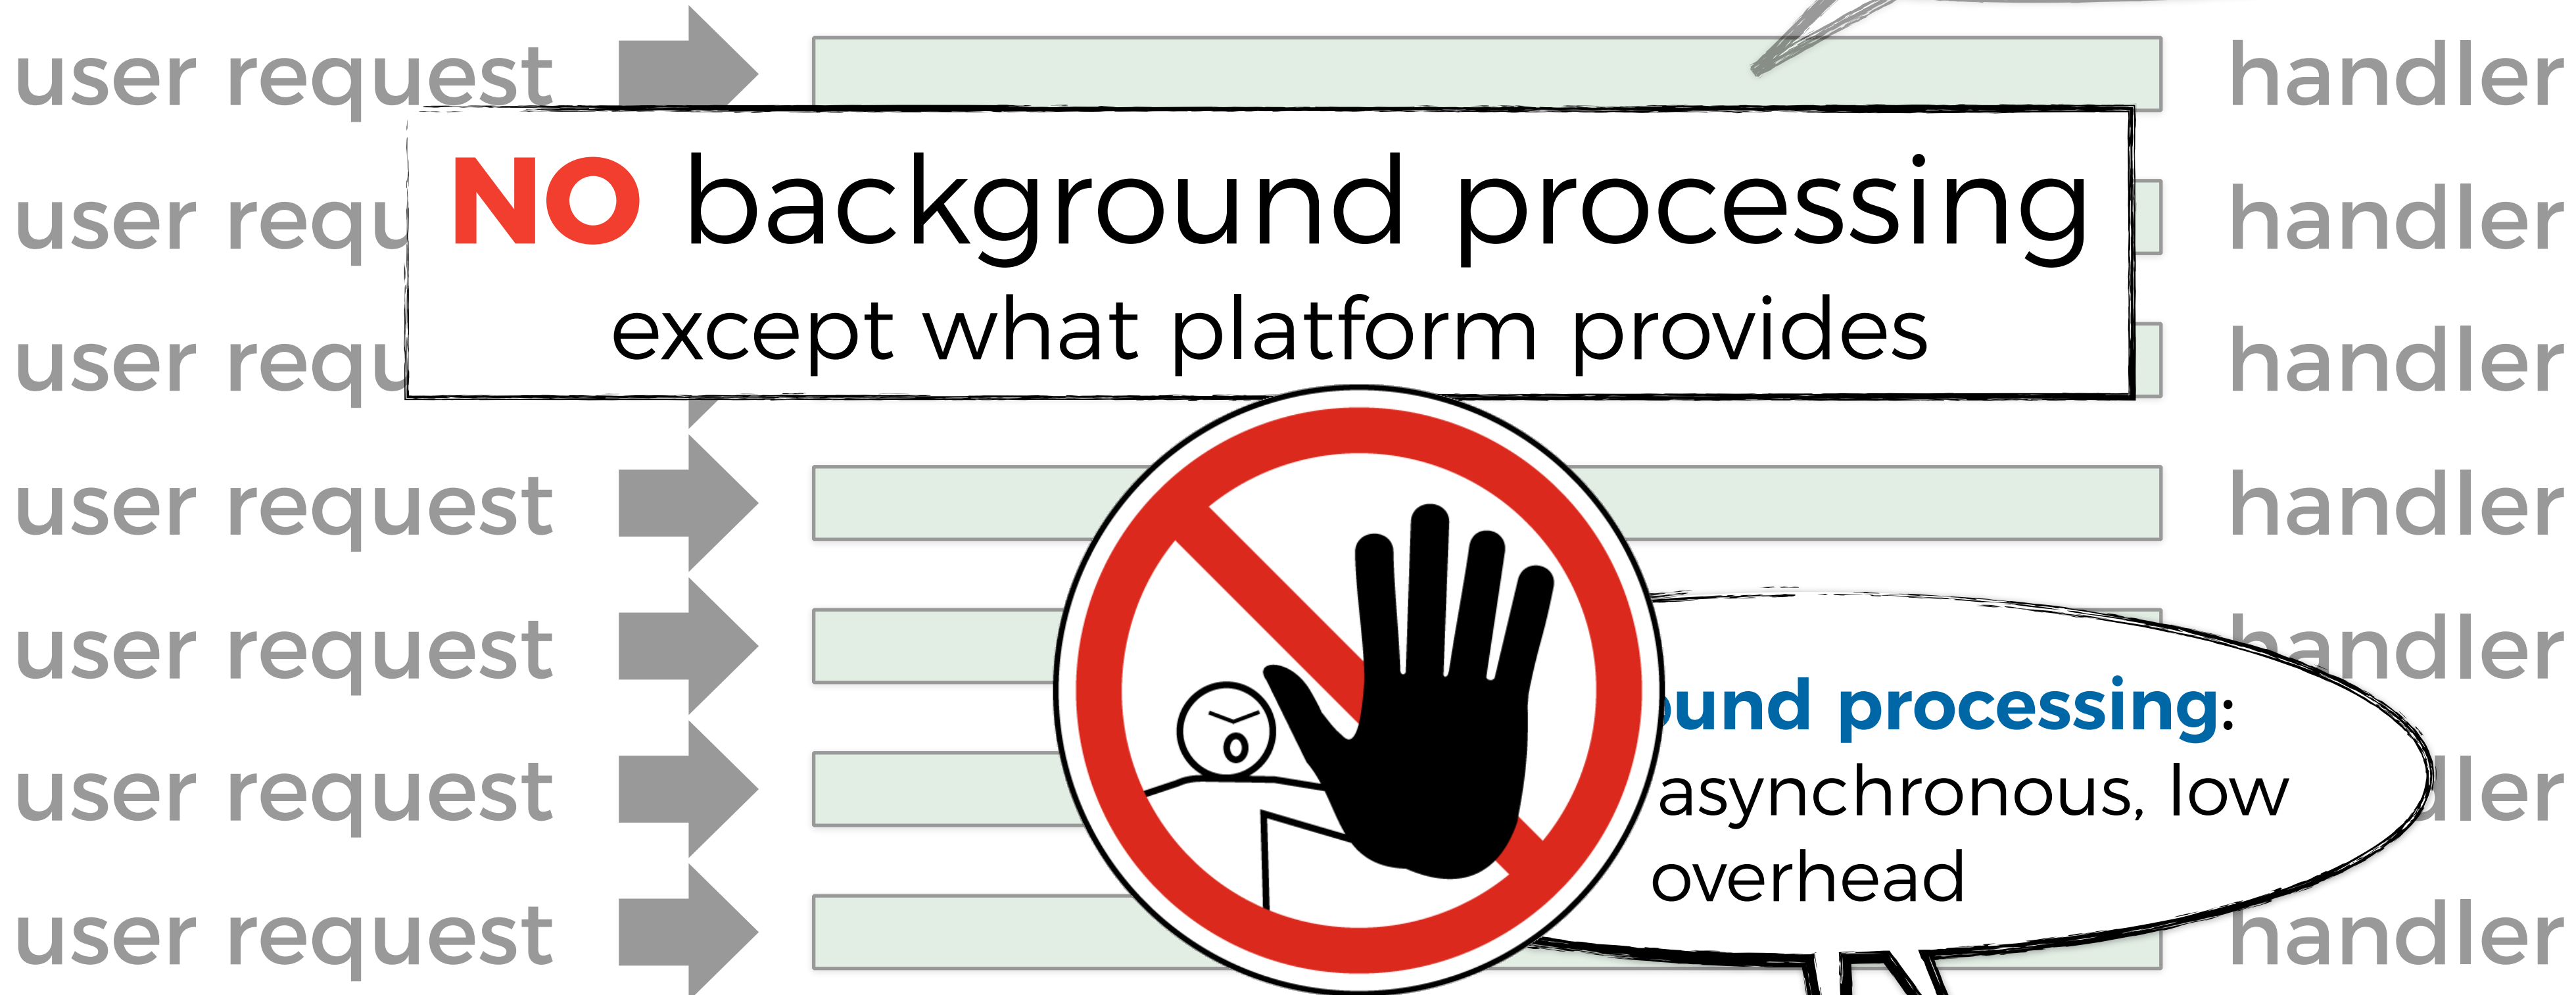
background processing:
batched, asynchronous, low overhead



StatsD

rsyslog

critical paths:
minimise user-facing latency

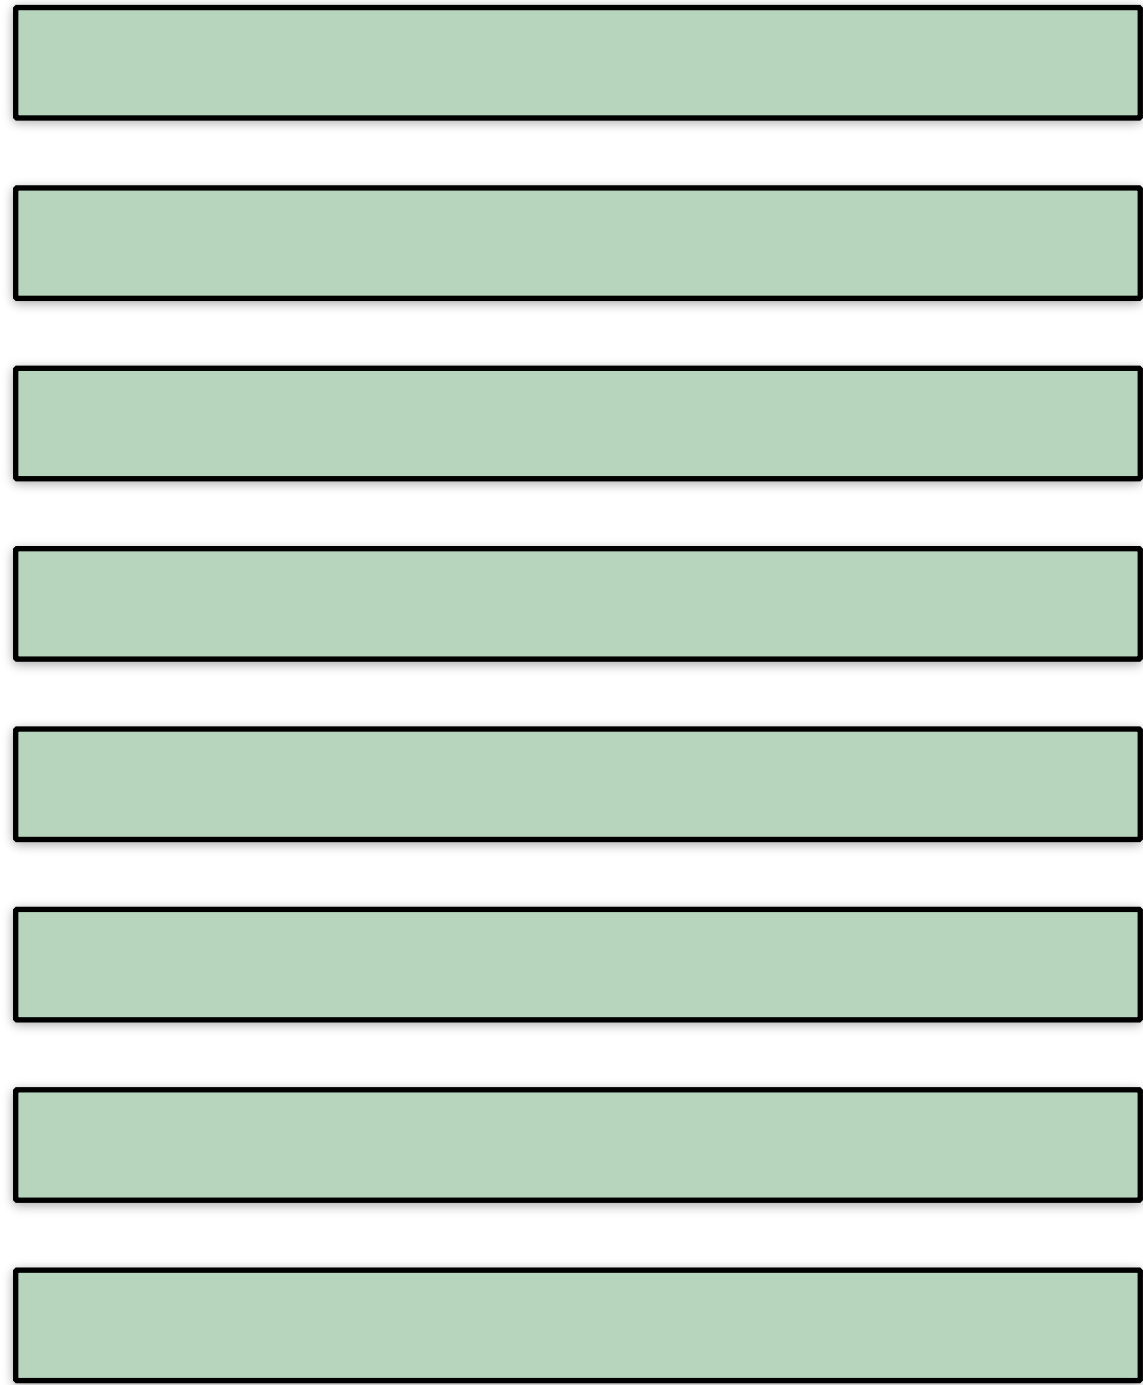


new challenges

- **nowhere** to install agents/daemons
- **no** background processing



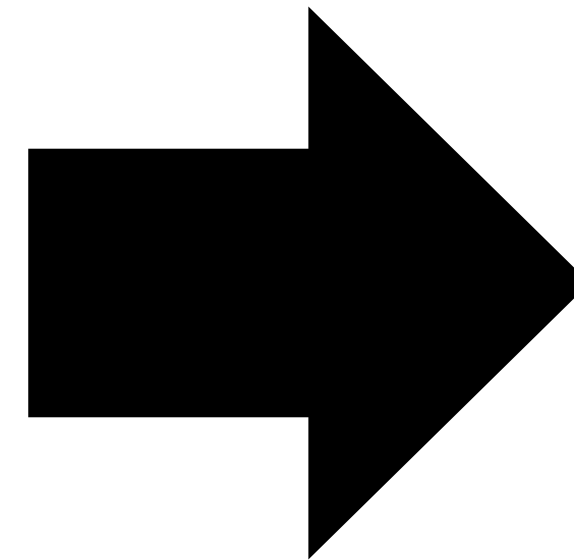
EC2



concurrency used to be handled by your code



EC2



Lambda



Lambda



Lambda



Lambda



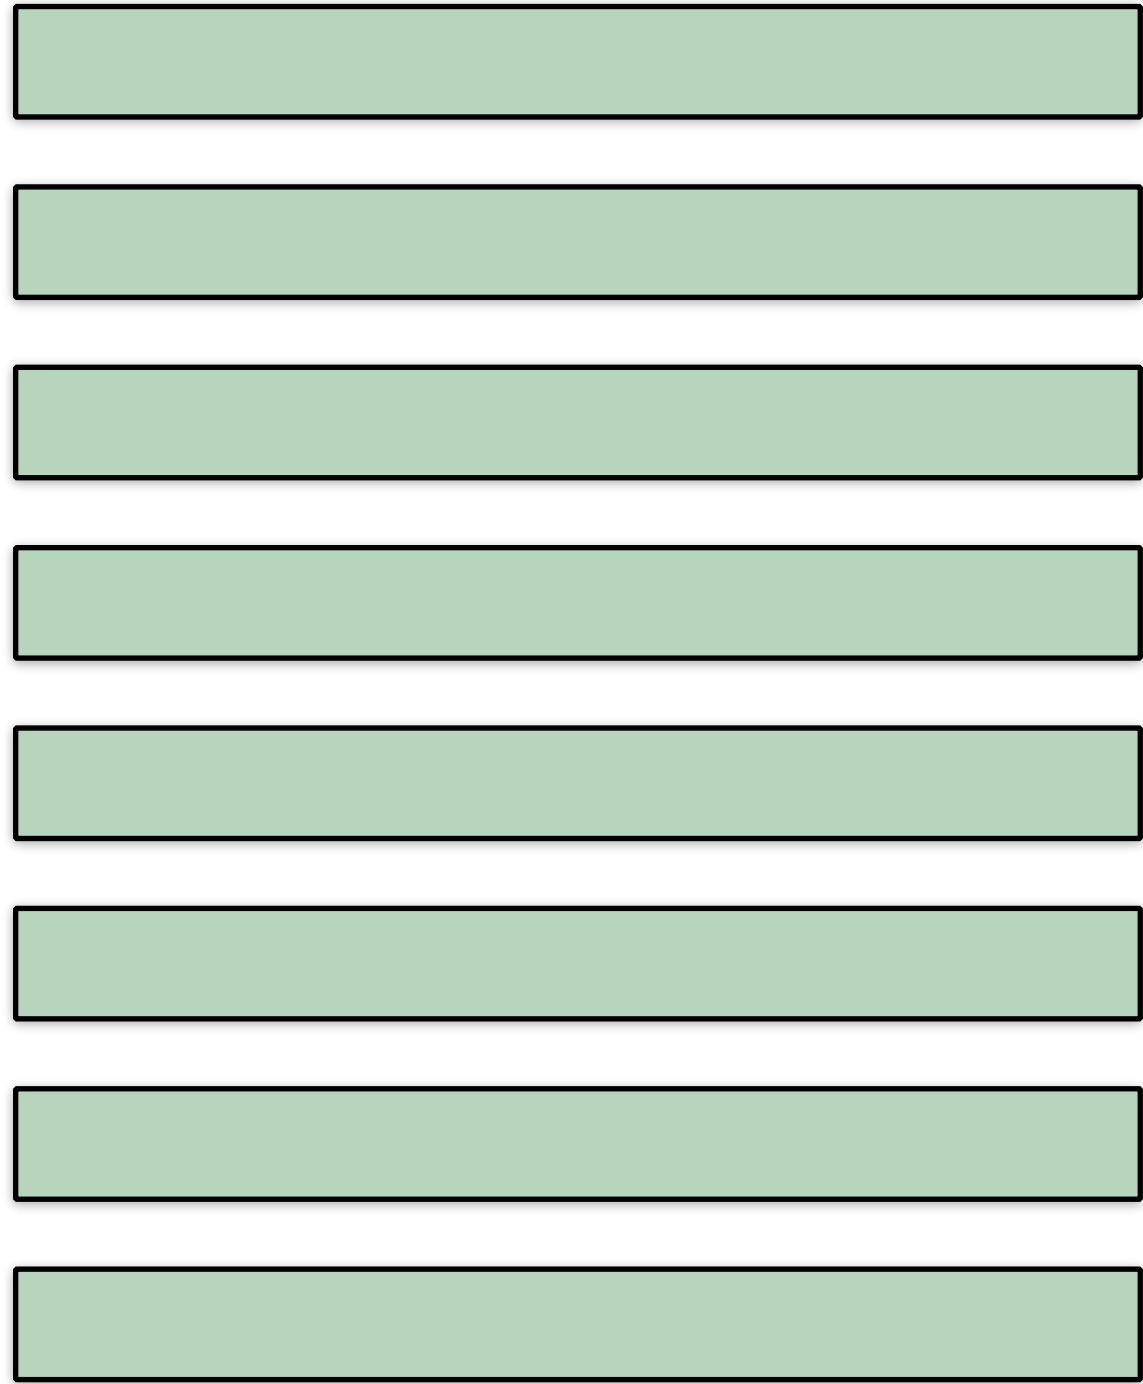
Lambda



now, it's handled by the AWS Lambda **platform**



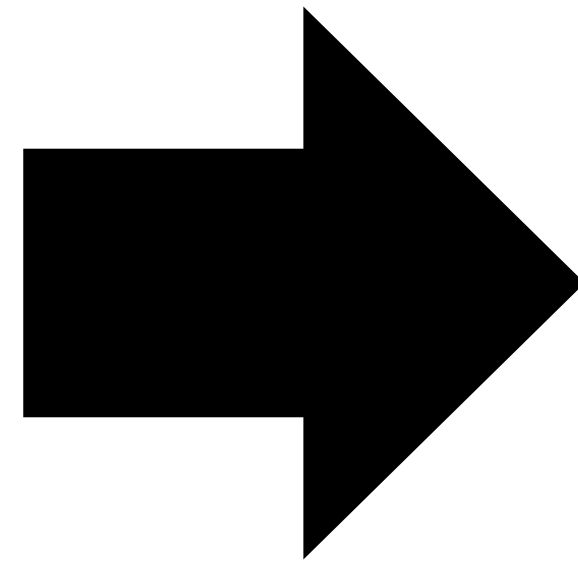
EC2



logs & metrics used to be **batched** here



EC2



Lambda



Lambda



Lambda



Lambda



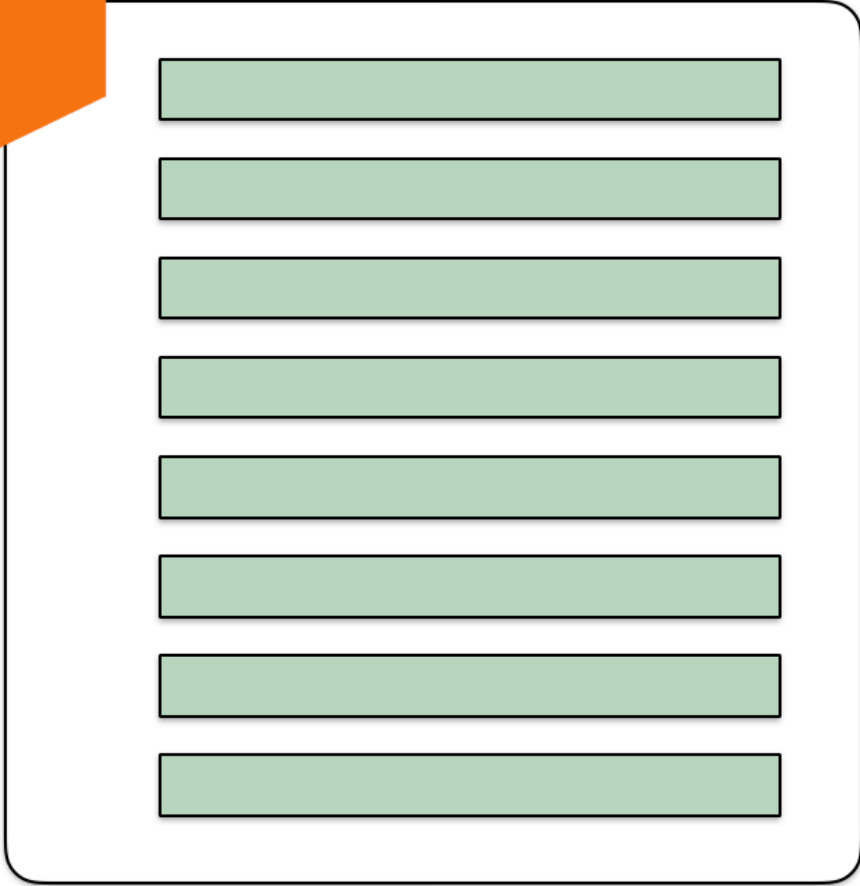
Lambda



now, they are batched in each **concurrent execution**, at best...

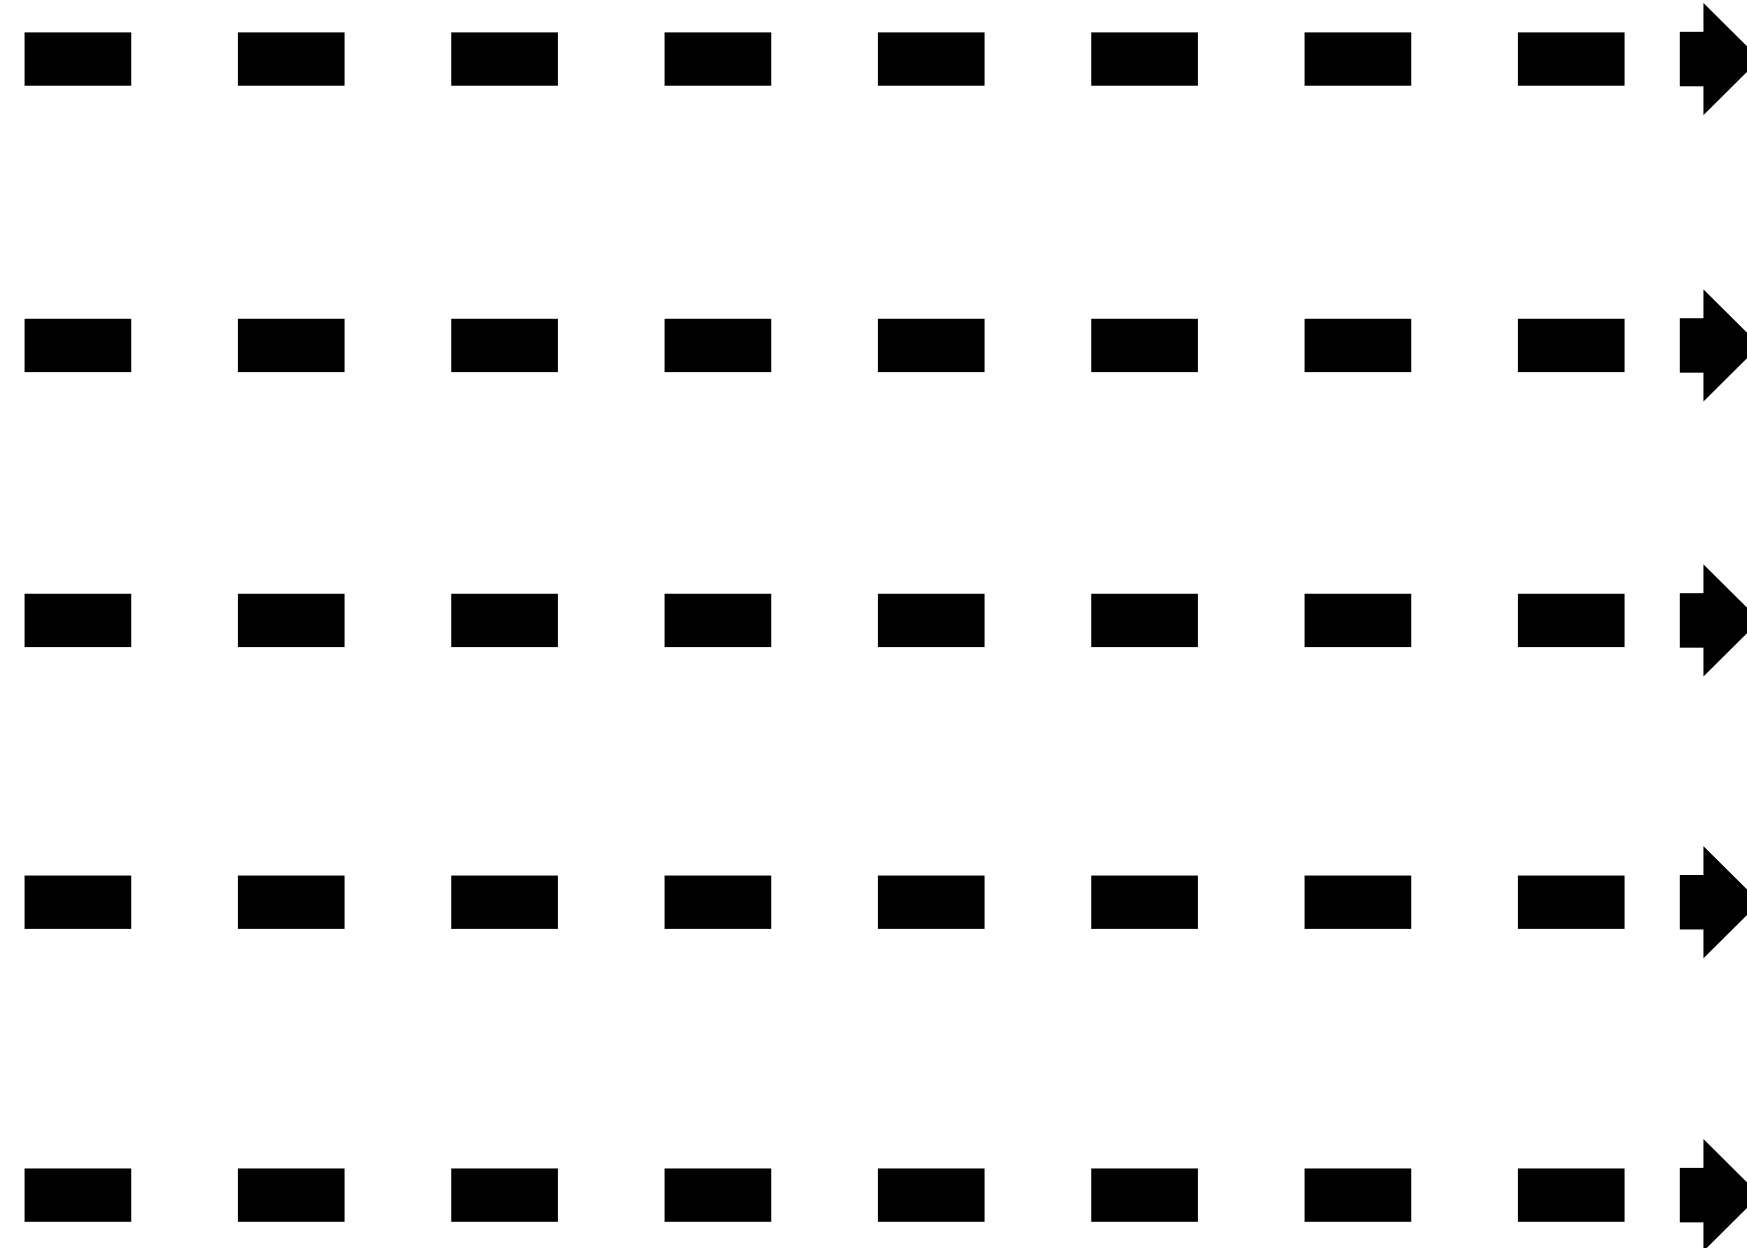
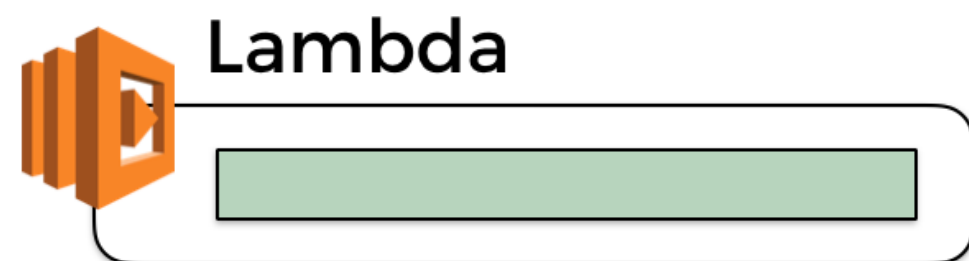
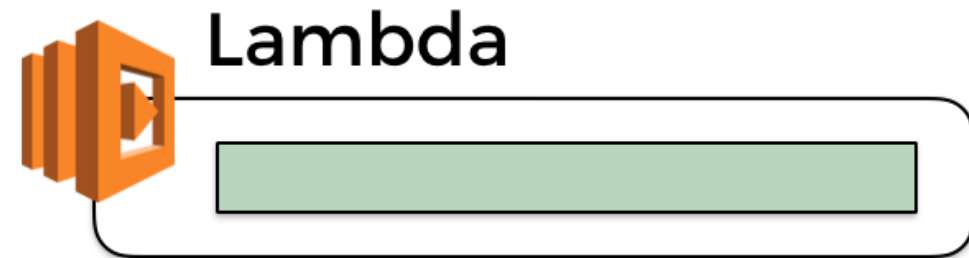
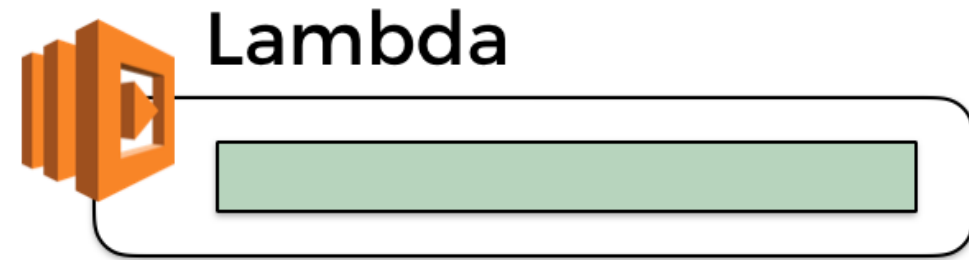
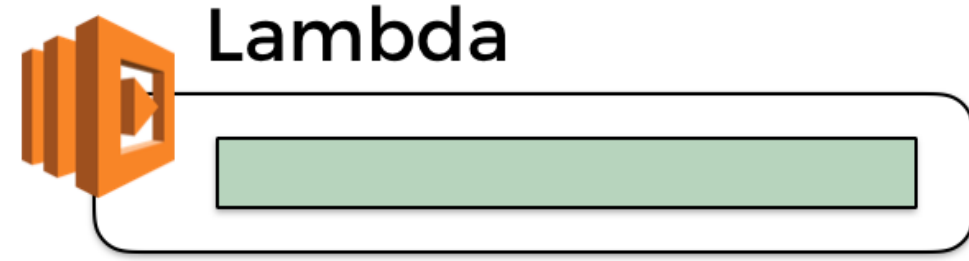
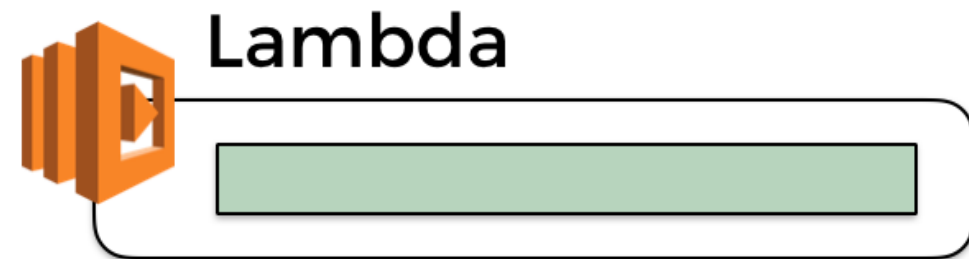
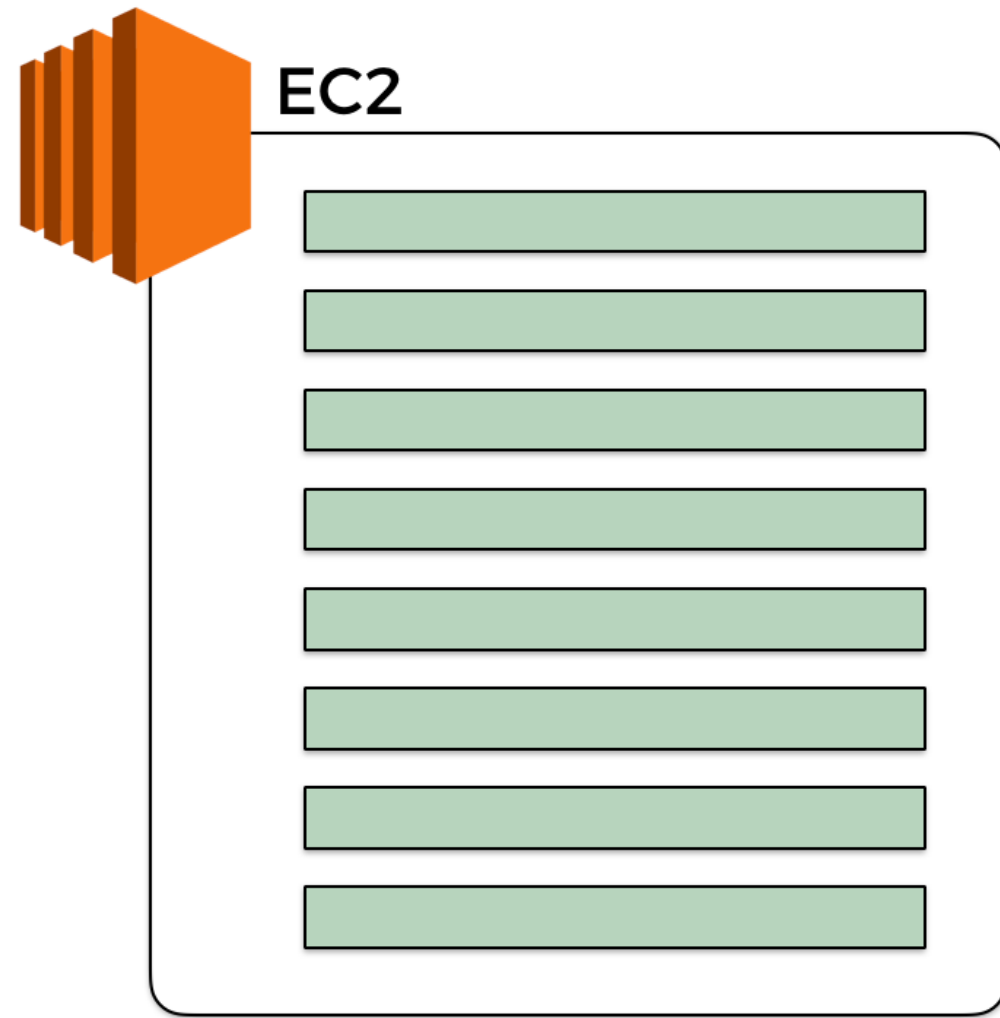


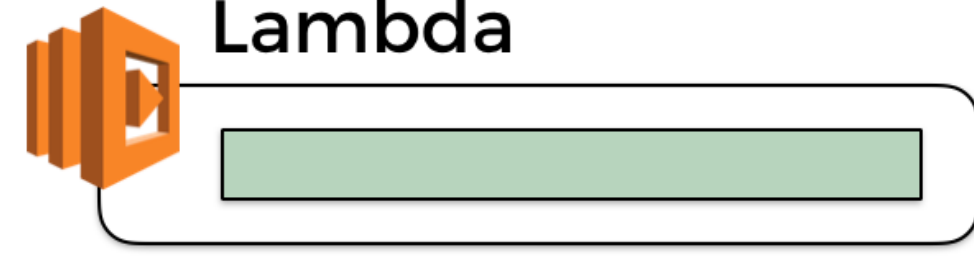
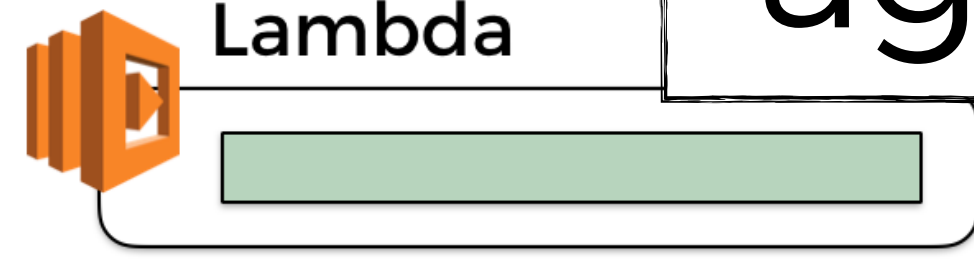
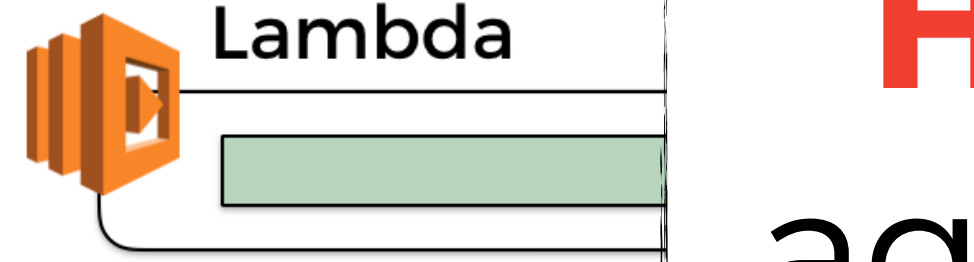
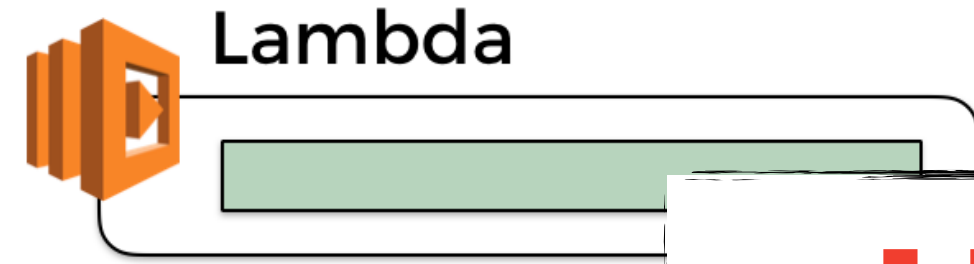
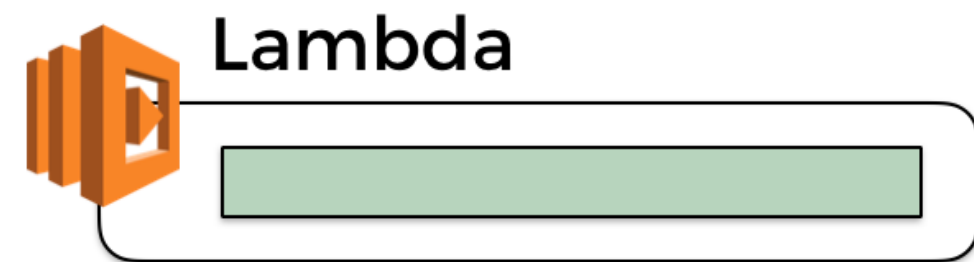
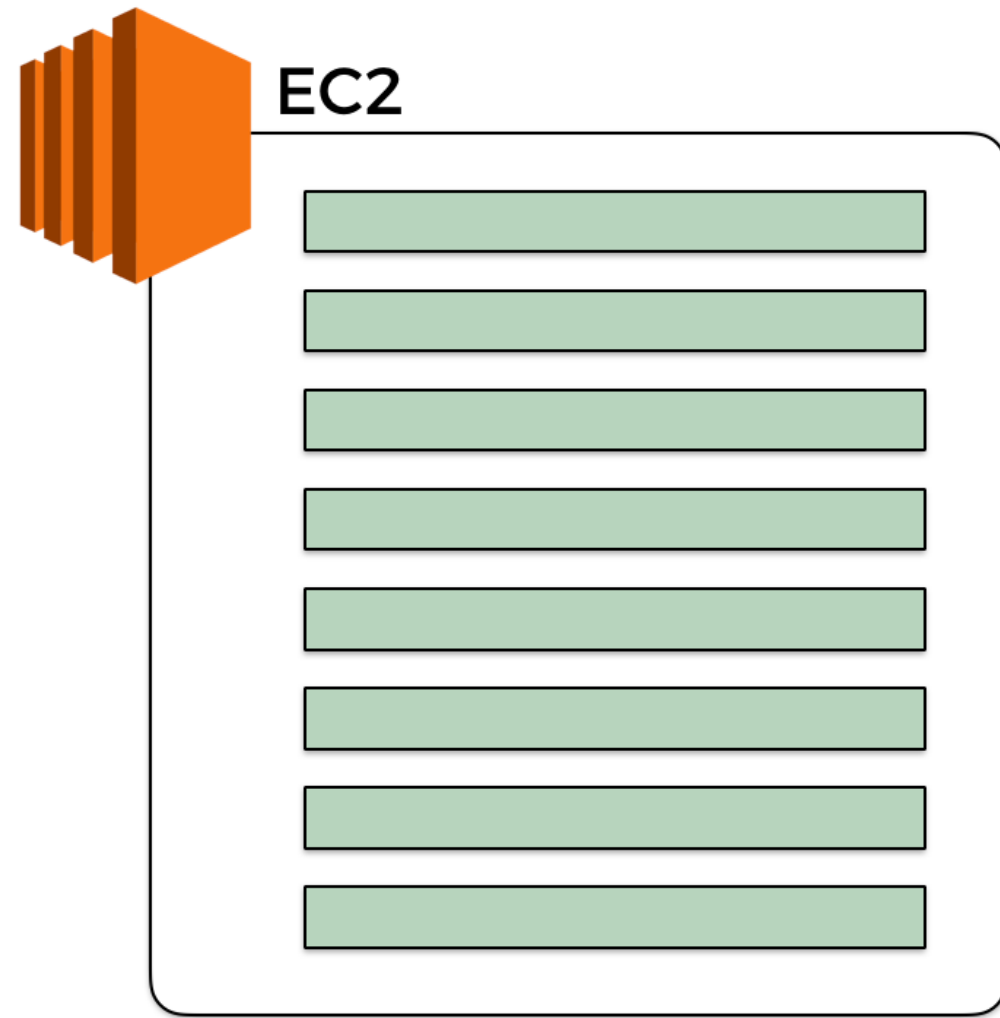
EC2



graphite







HIGHER concurrency to log aggregation/telemetry system



new challenges

- **nowhere** to install agents/daemons
- **no** background processing
- **higher** concurrency to telemetry system



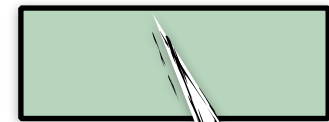
Lambda



cold start



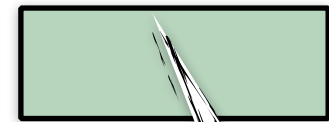
Lambda



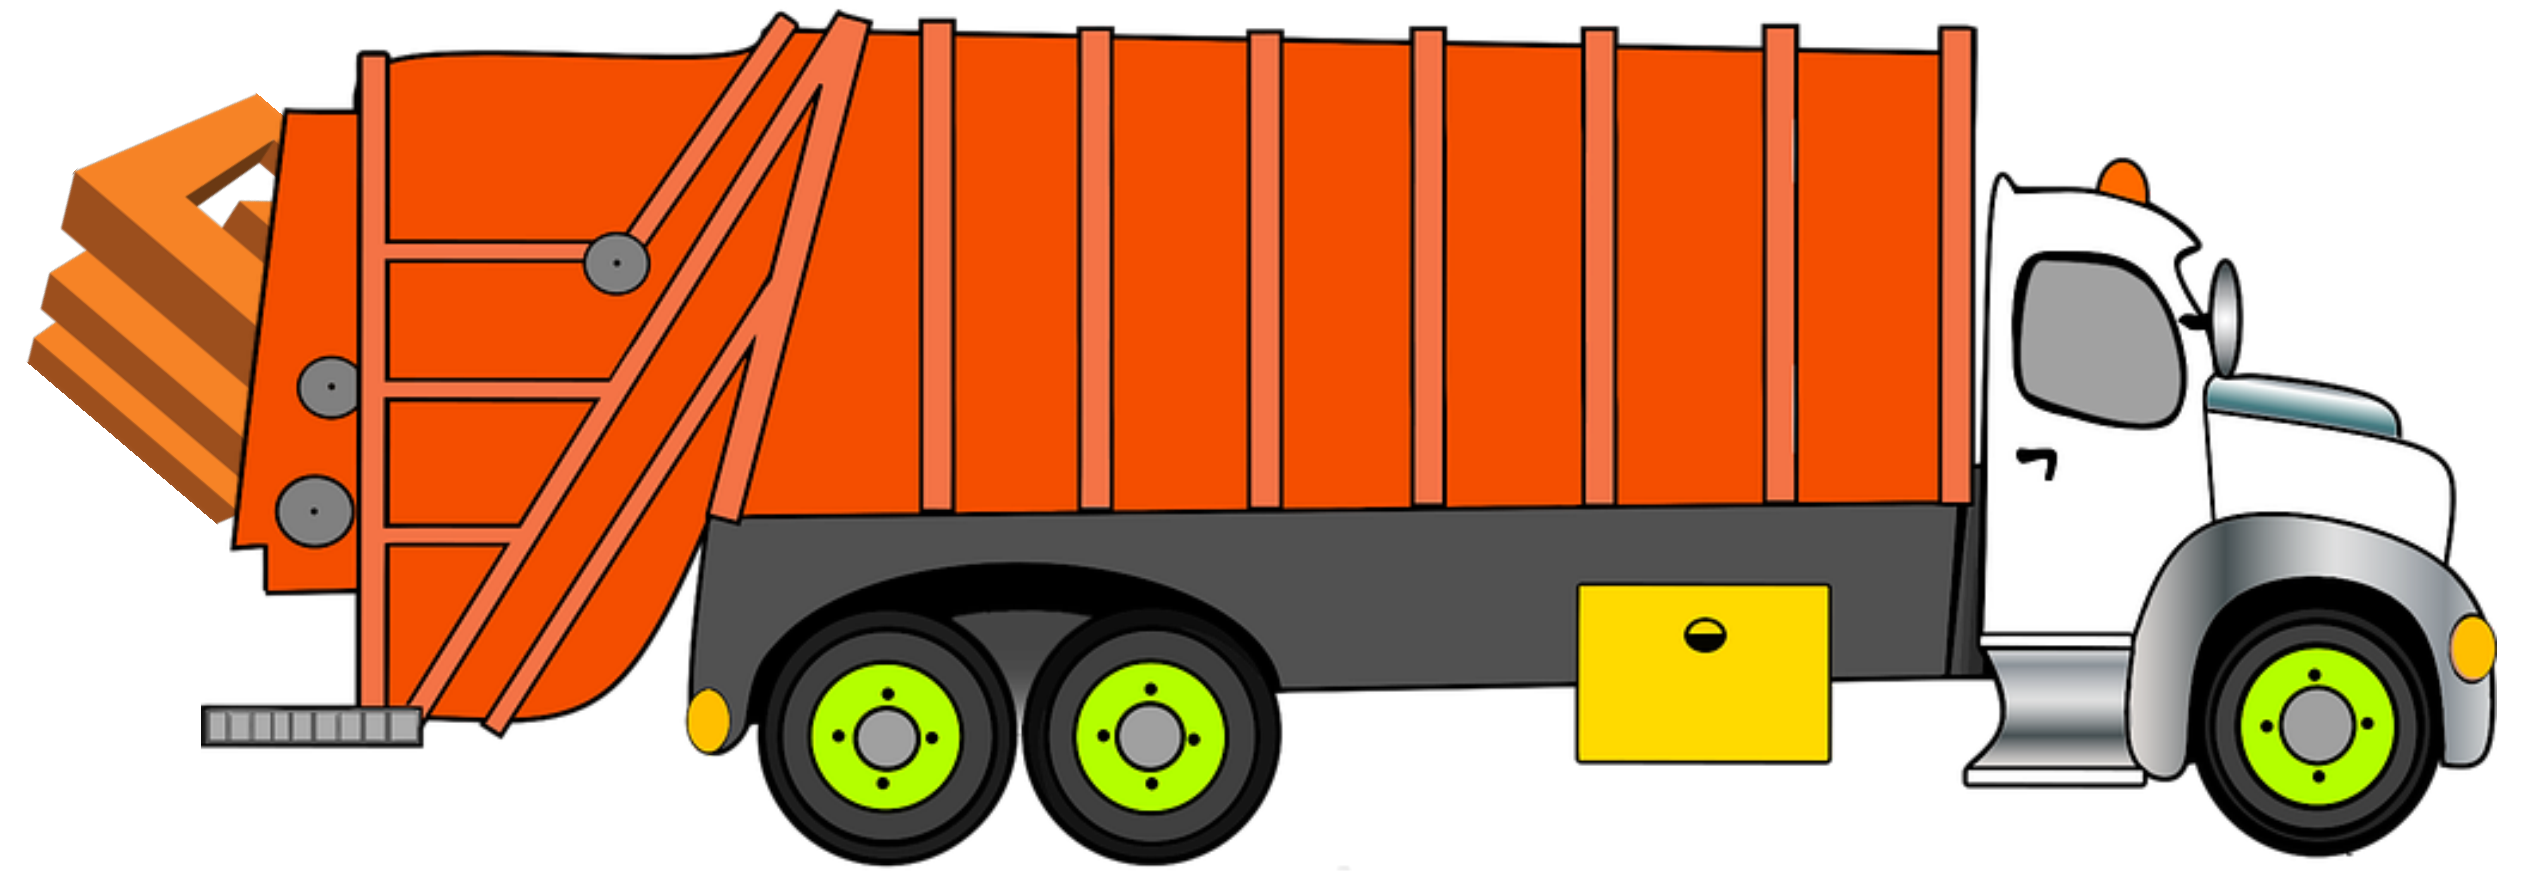
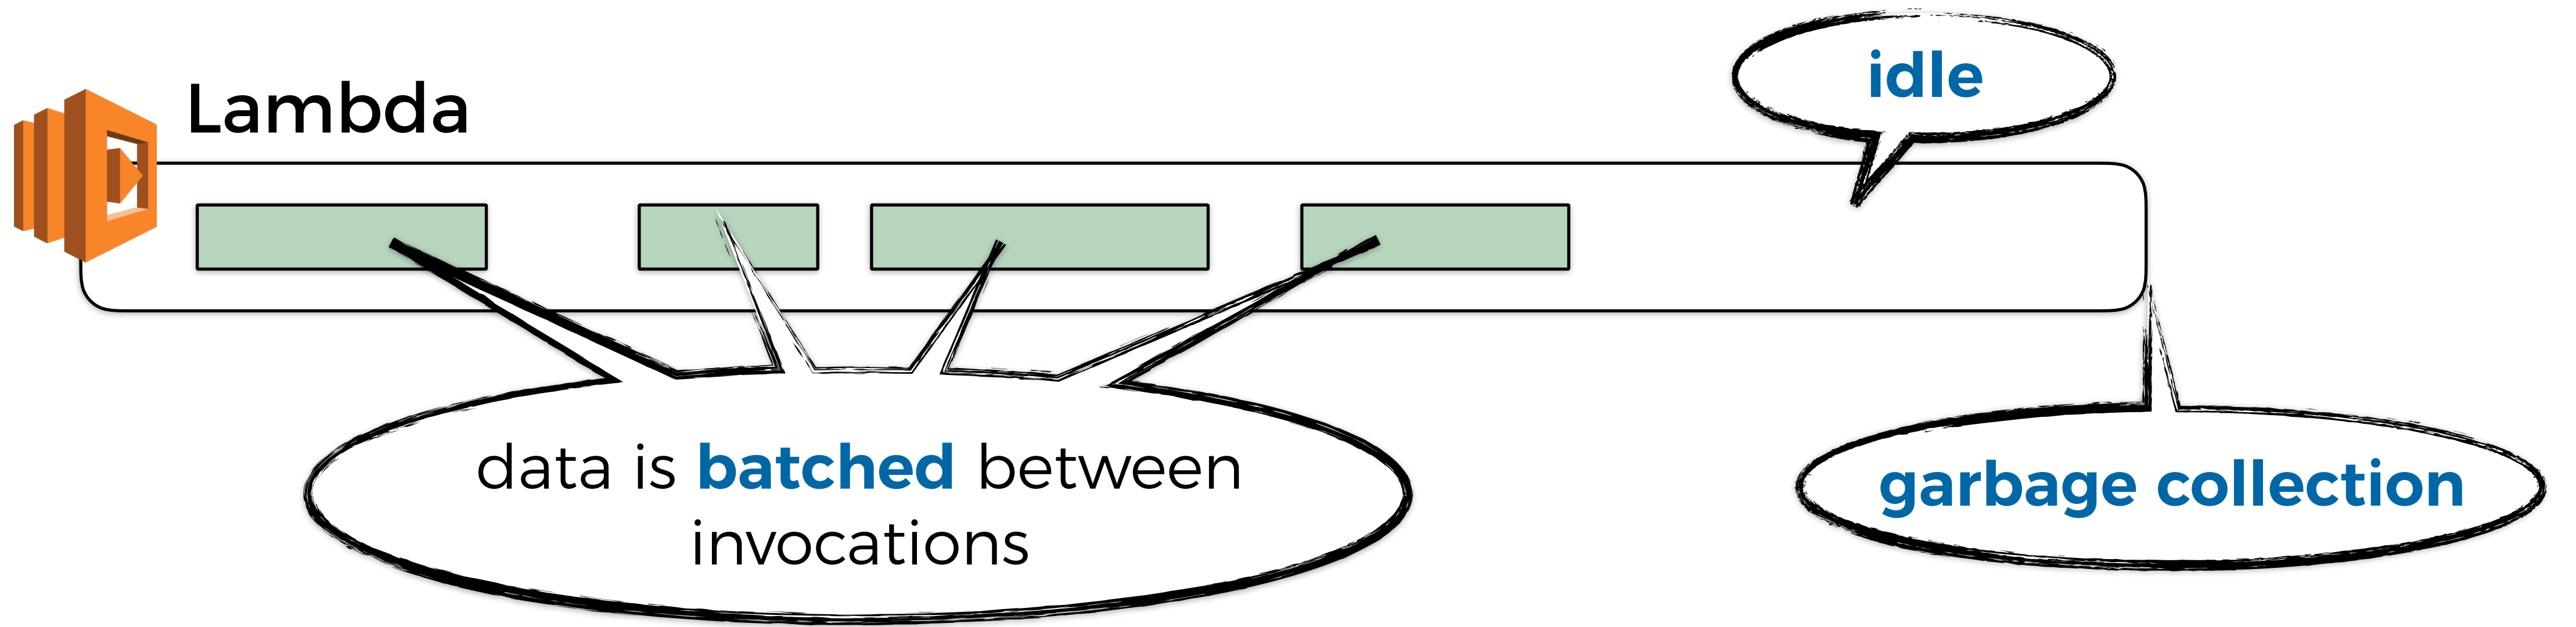
data is **batched** between
invocations

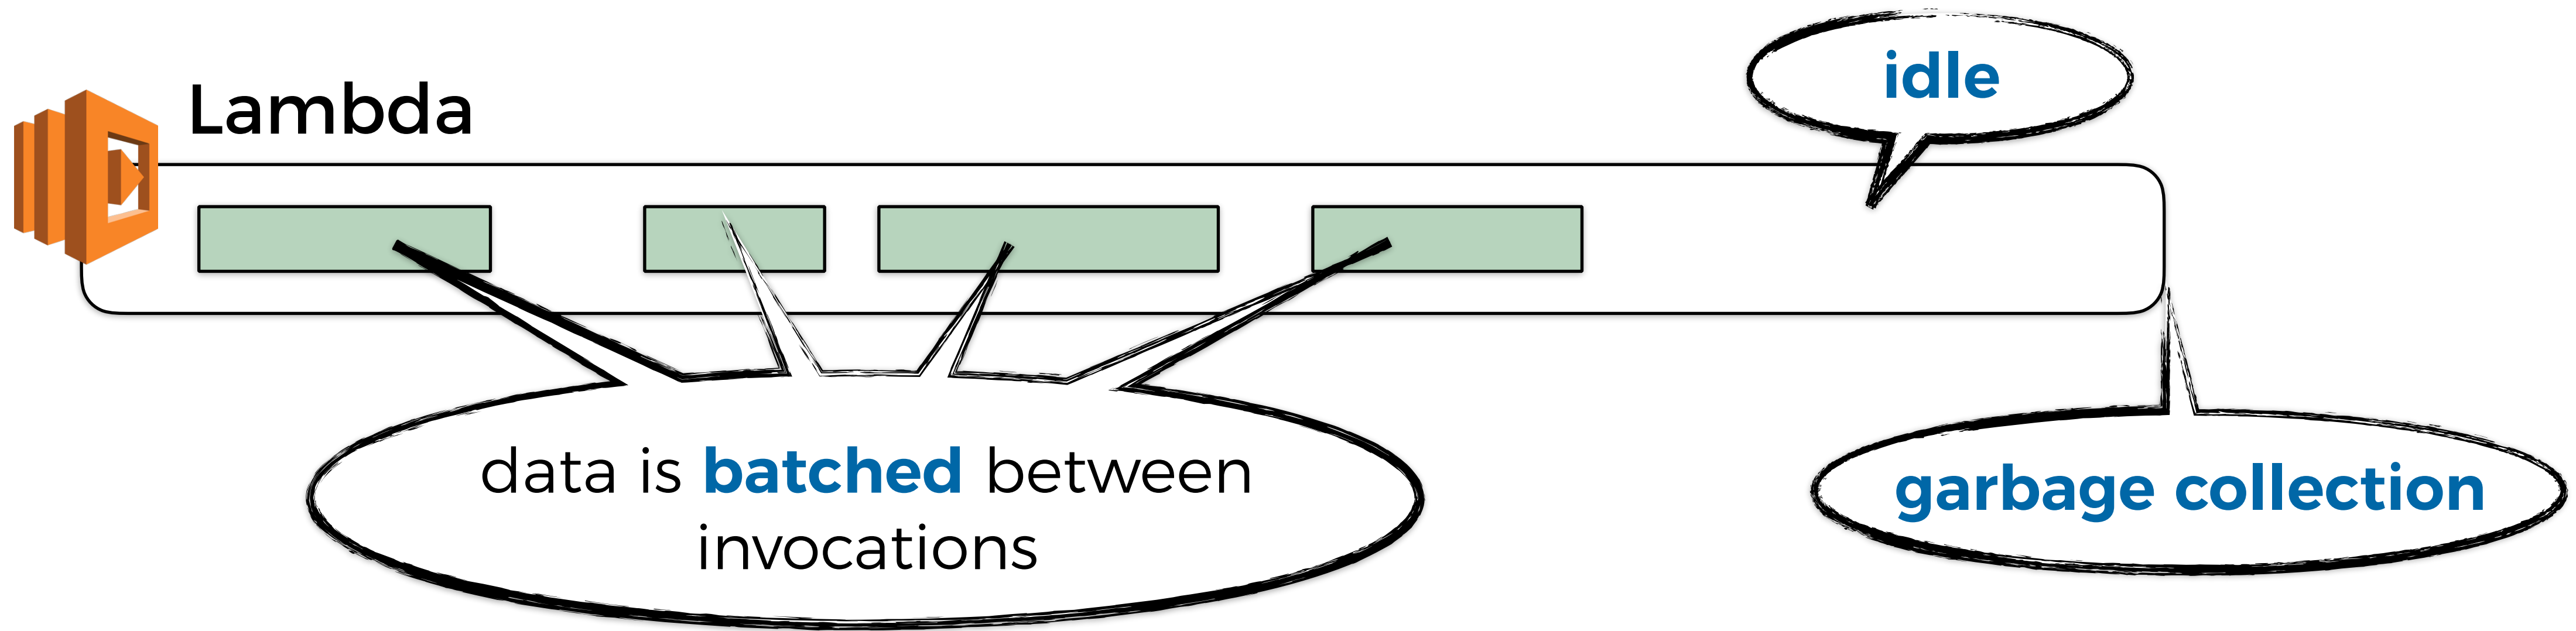


Lambda



data is **batched** between
invocations





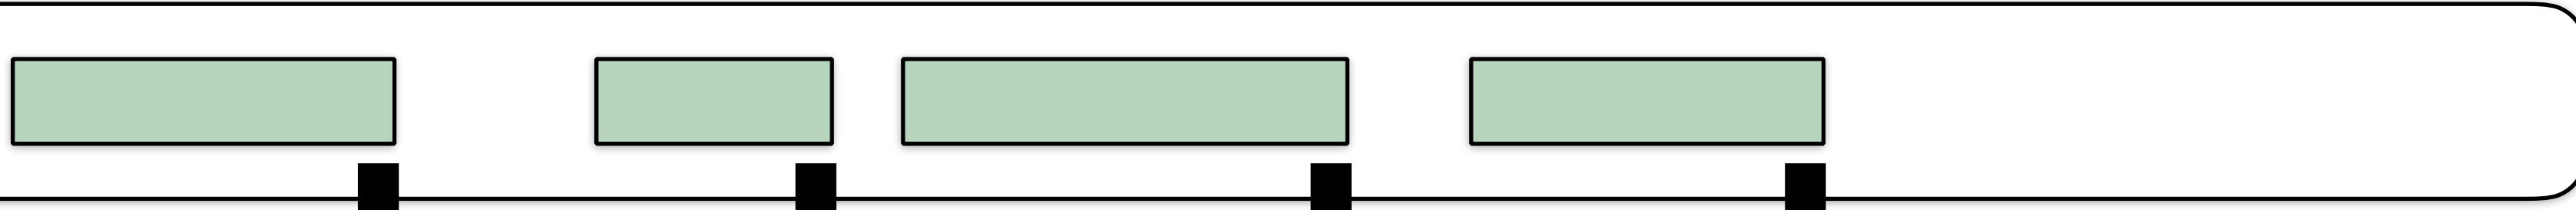
HIGH chance of data loss

new challenges

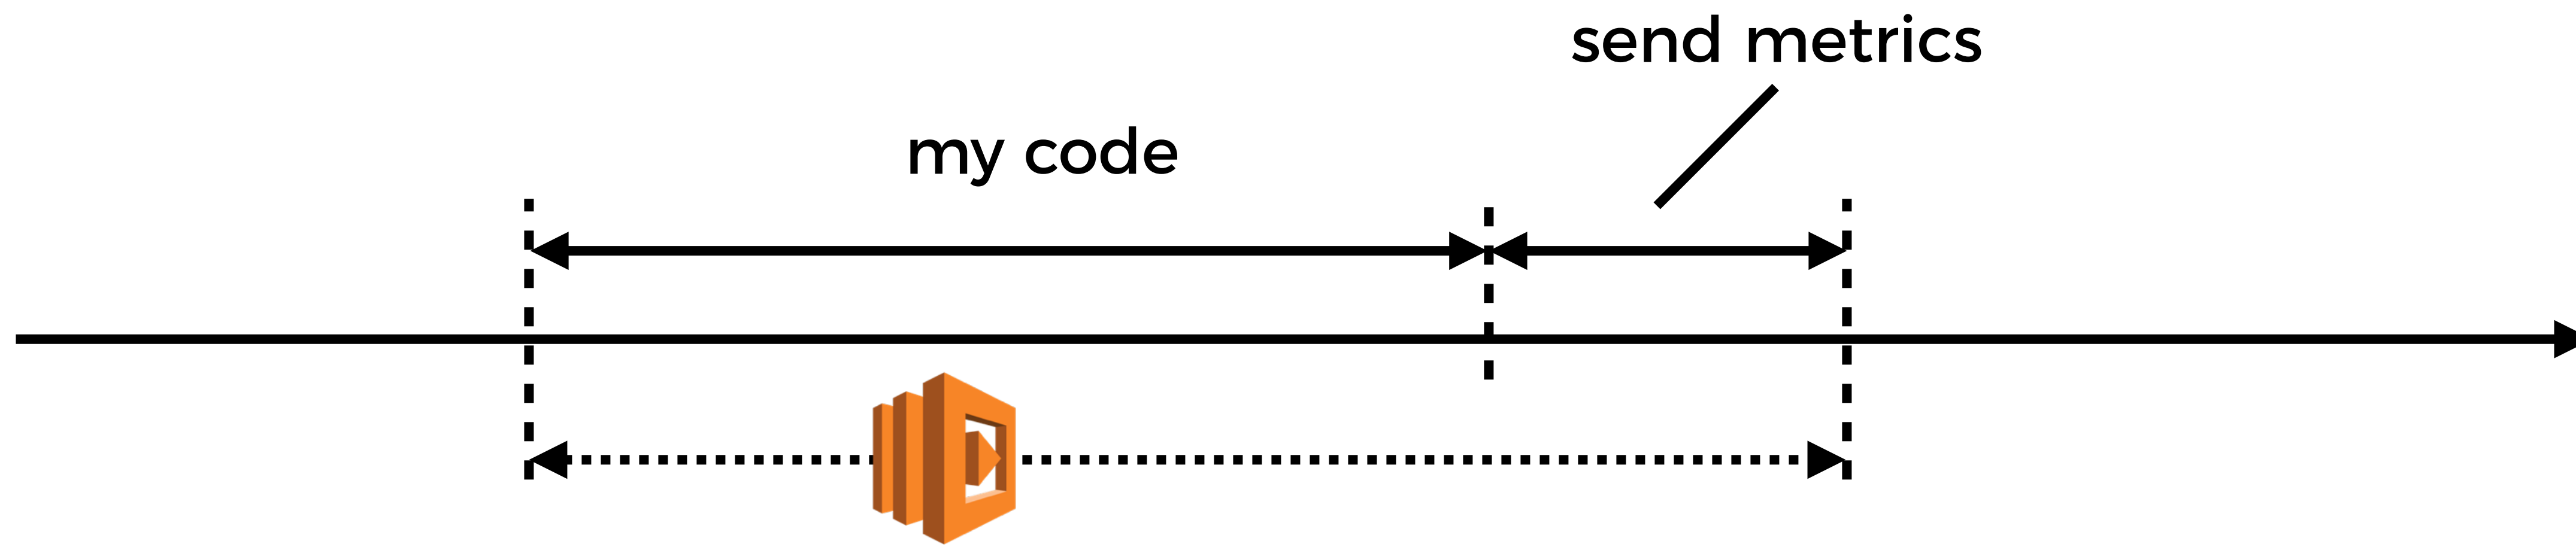
- **nowhere** to install agents/daemons
- **no** background processing
- **higher** concurrency to telemetry system
- **high** chance of data loss (if batching)

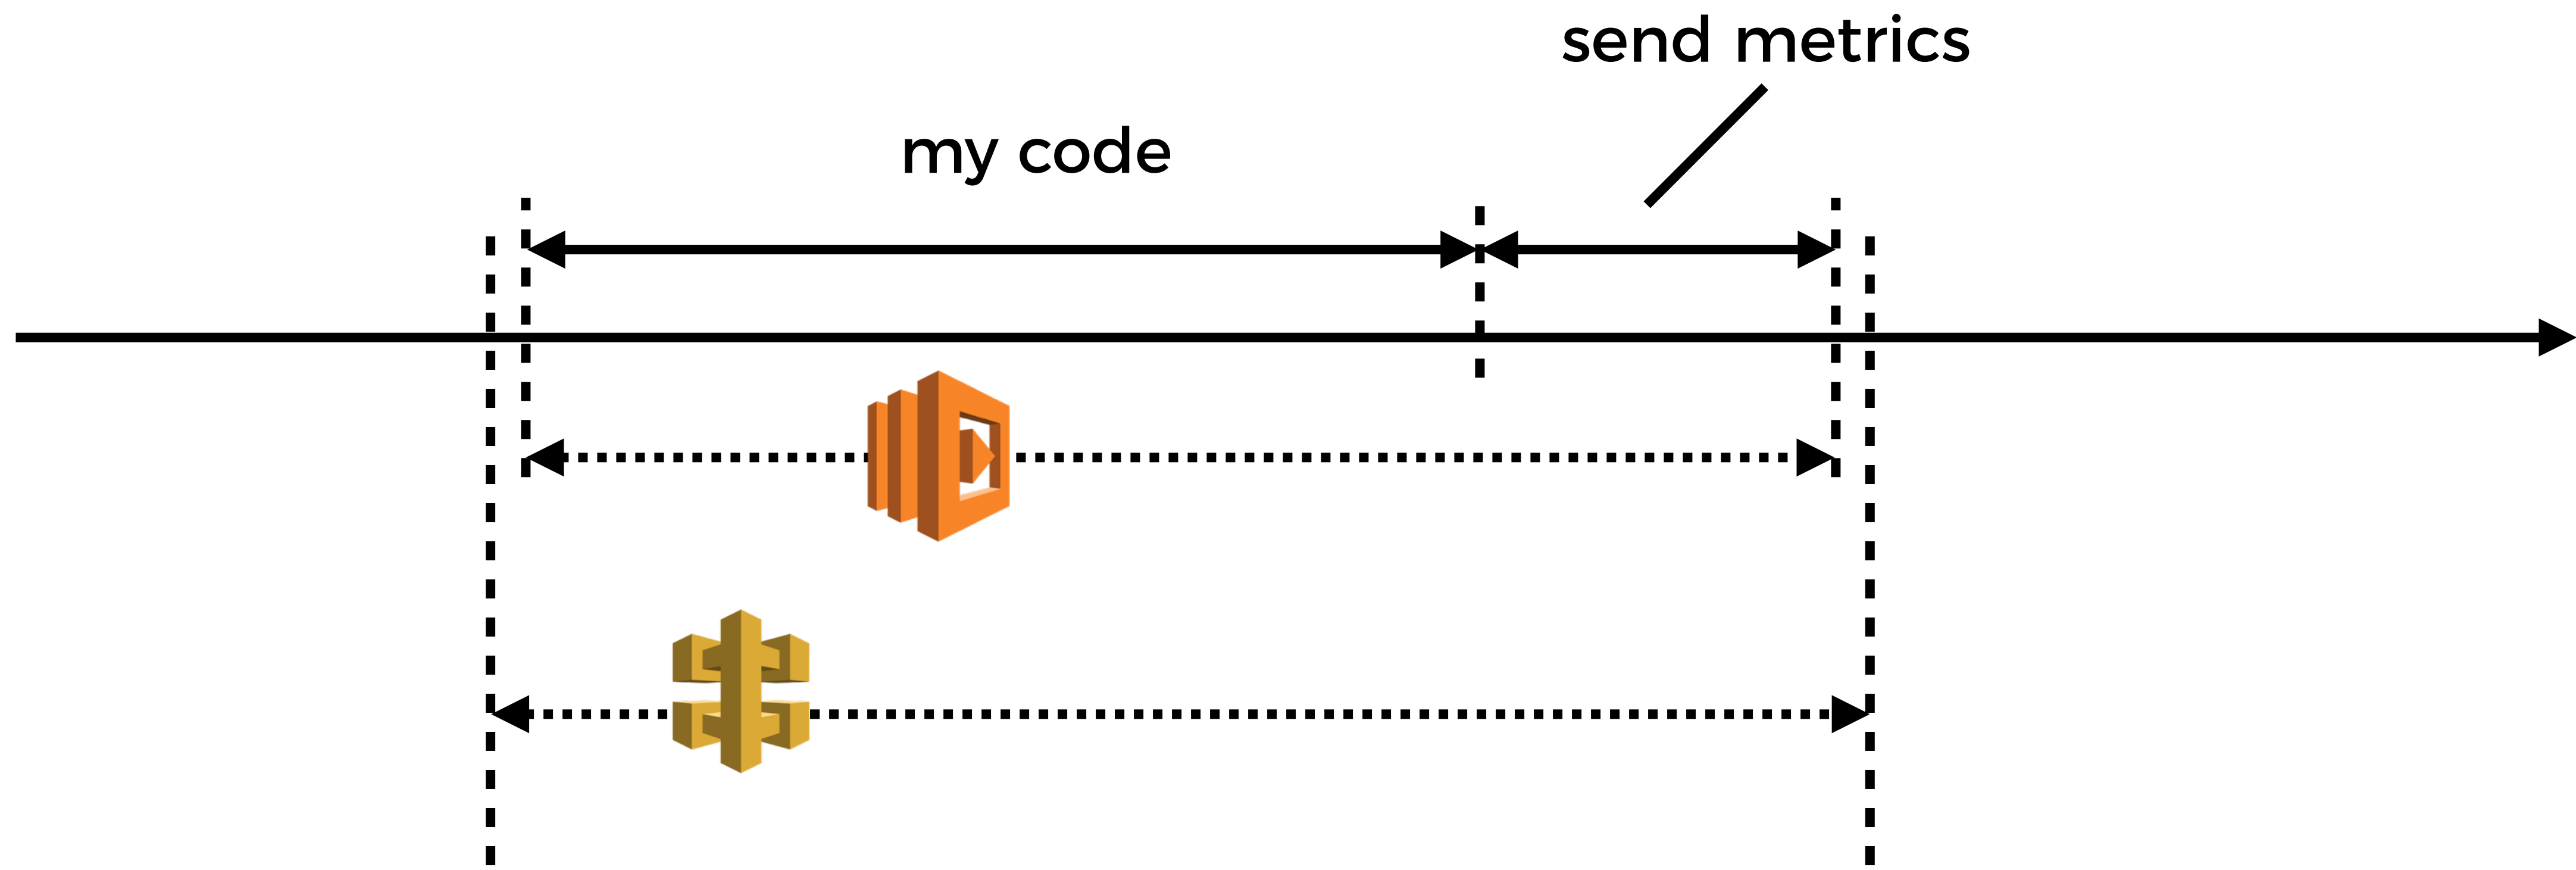


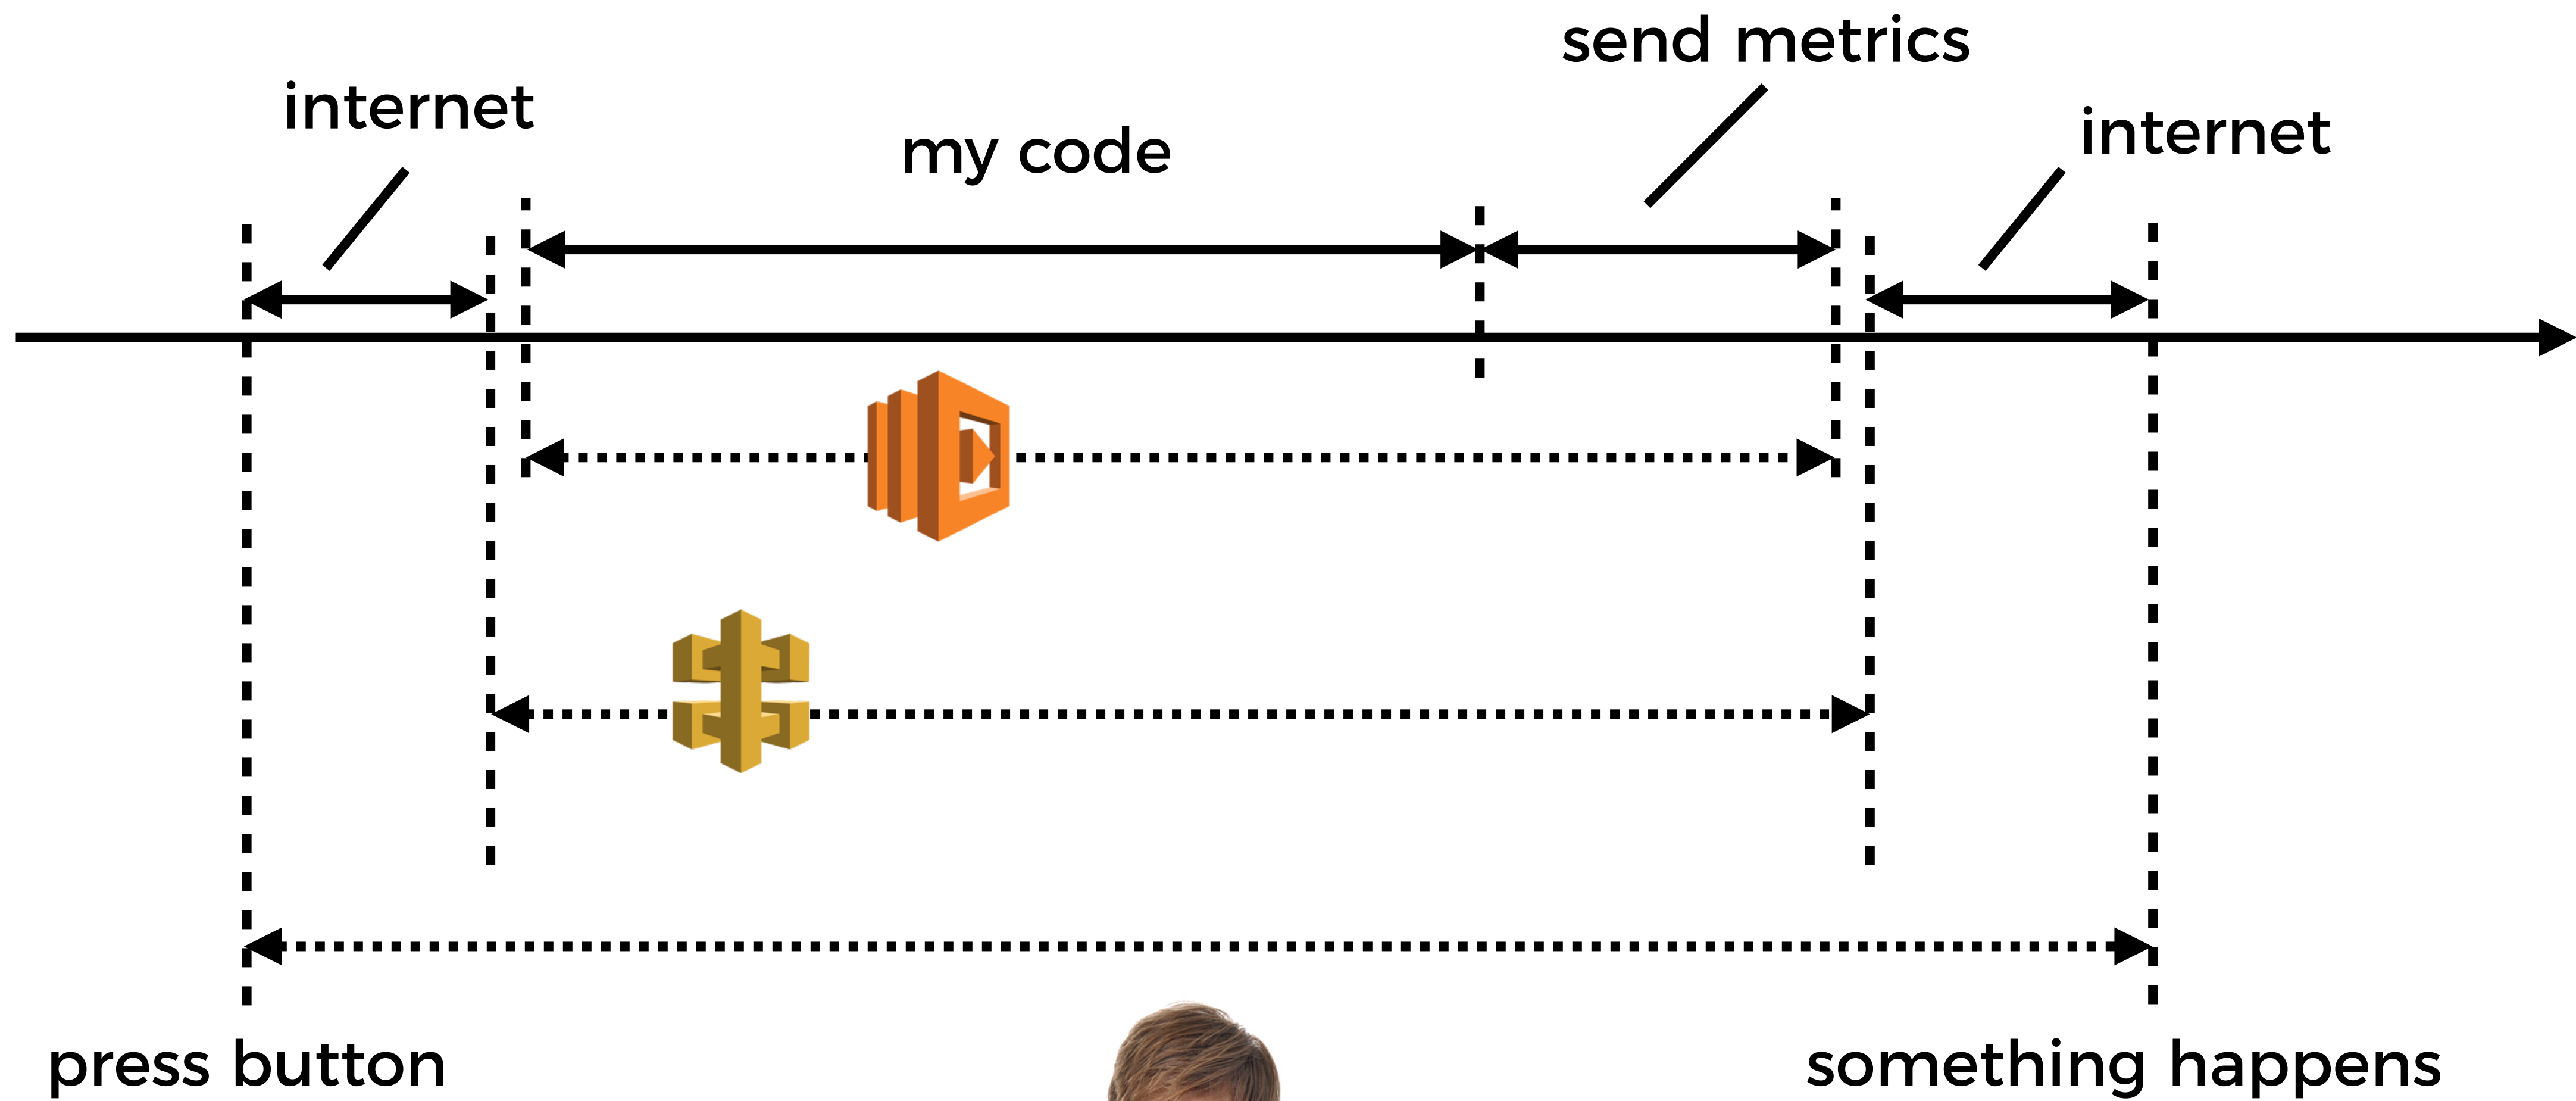
Lambda



graphite











Randy Shoup [Follow](#)
Jan 5 · 8 min read



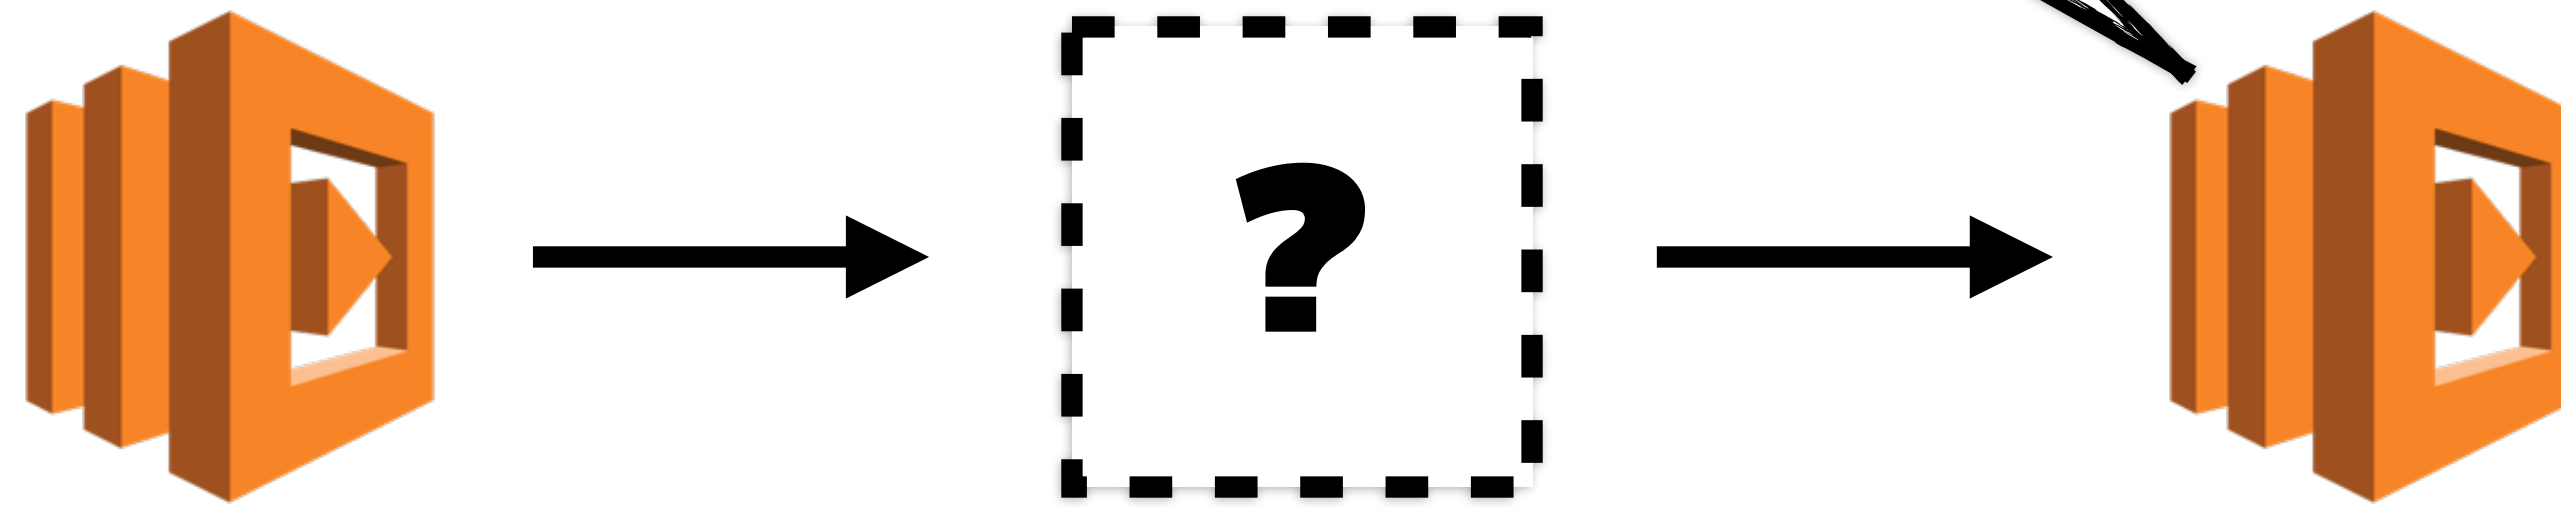
[Image Source](#)

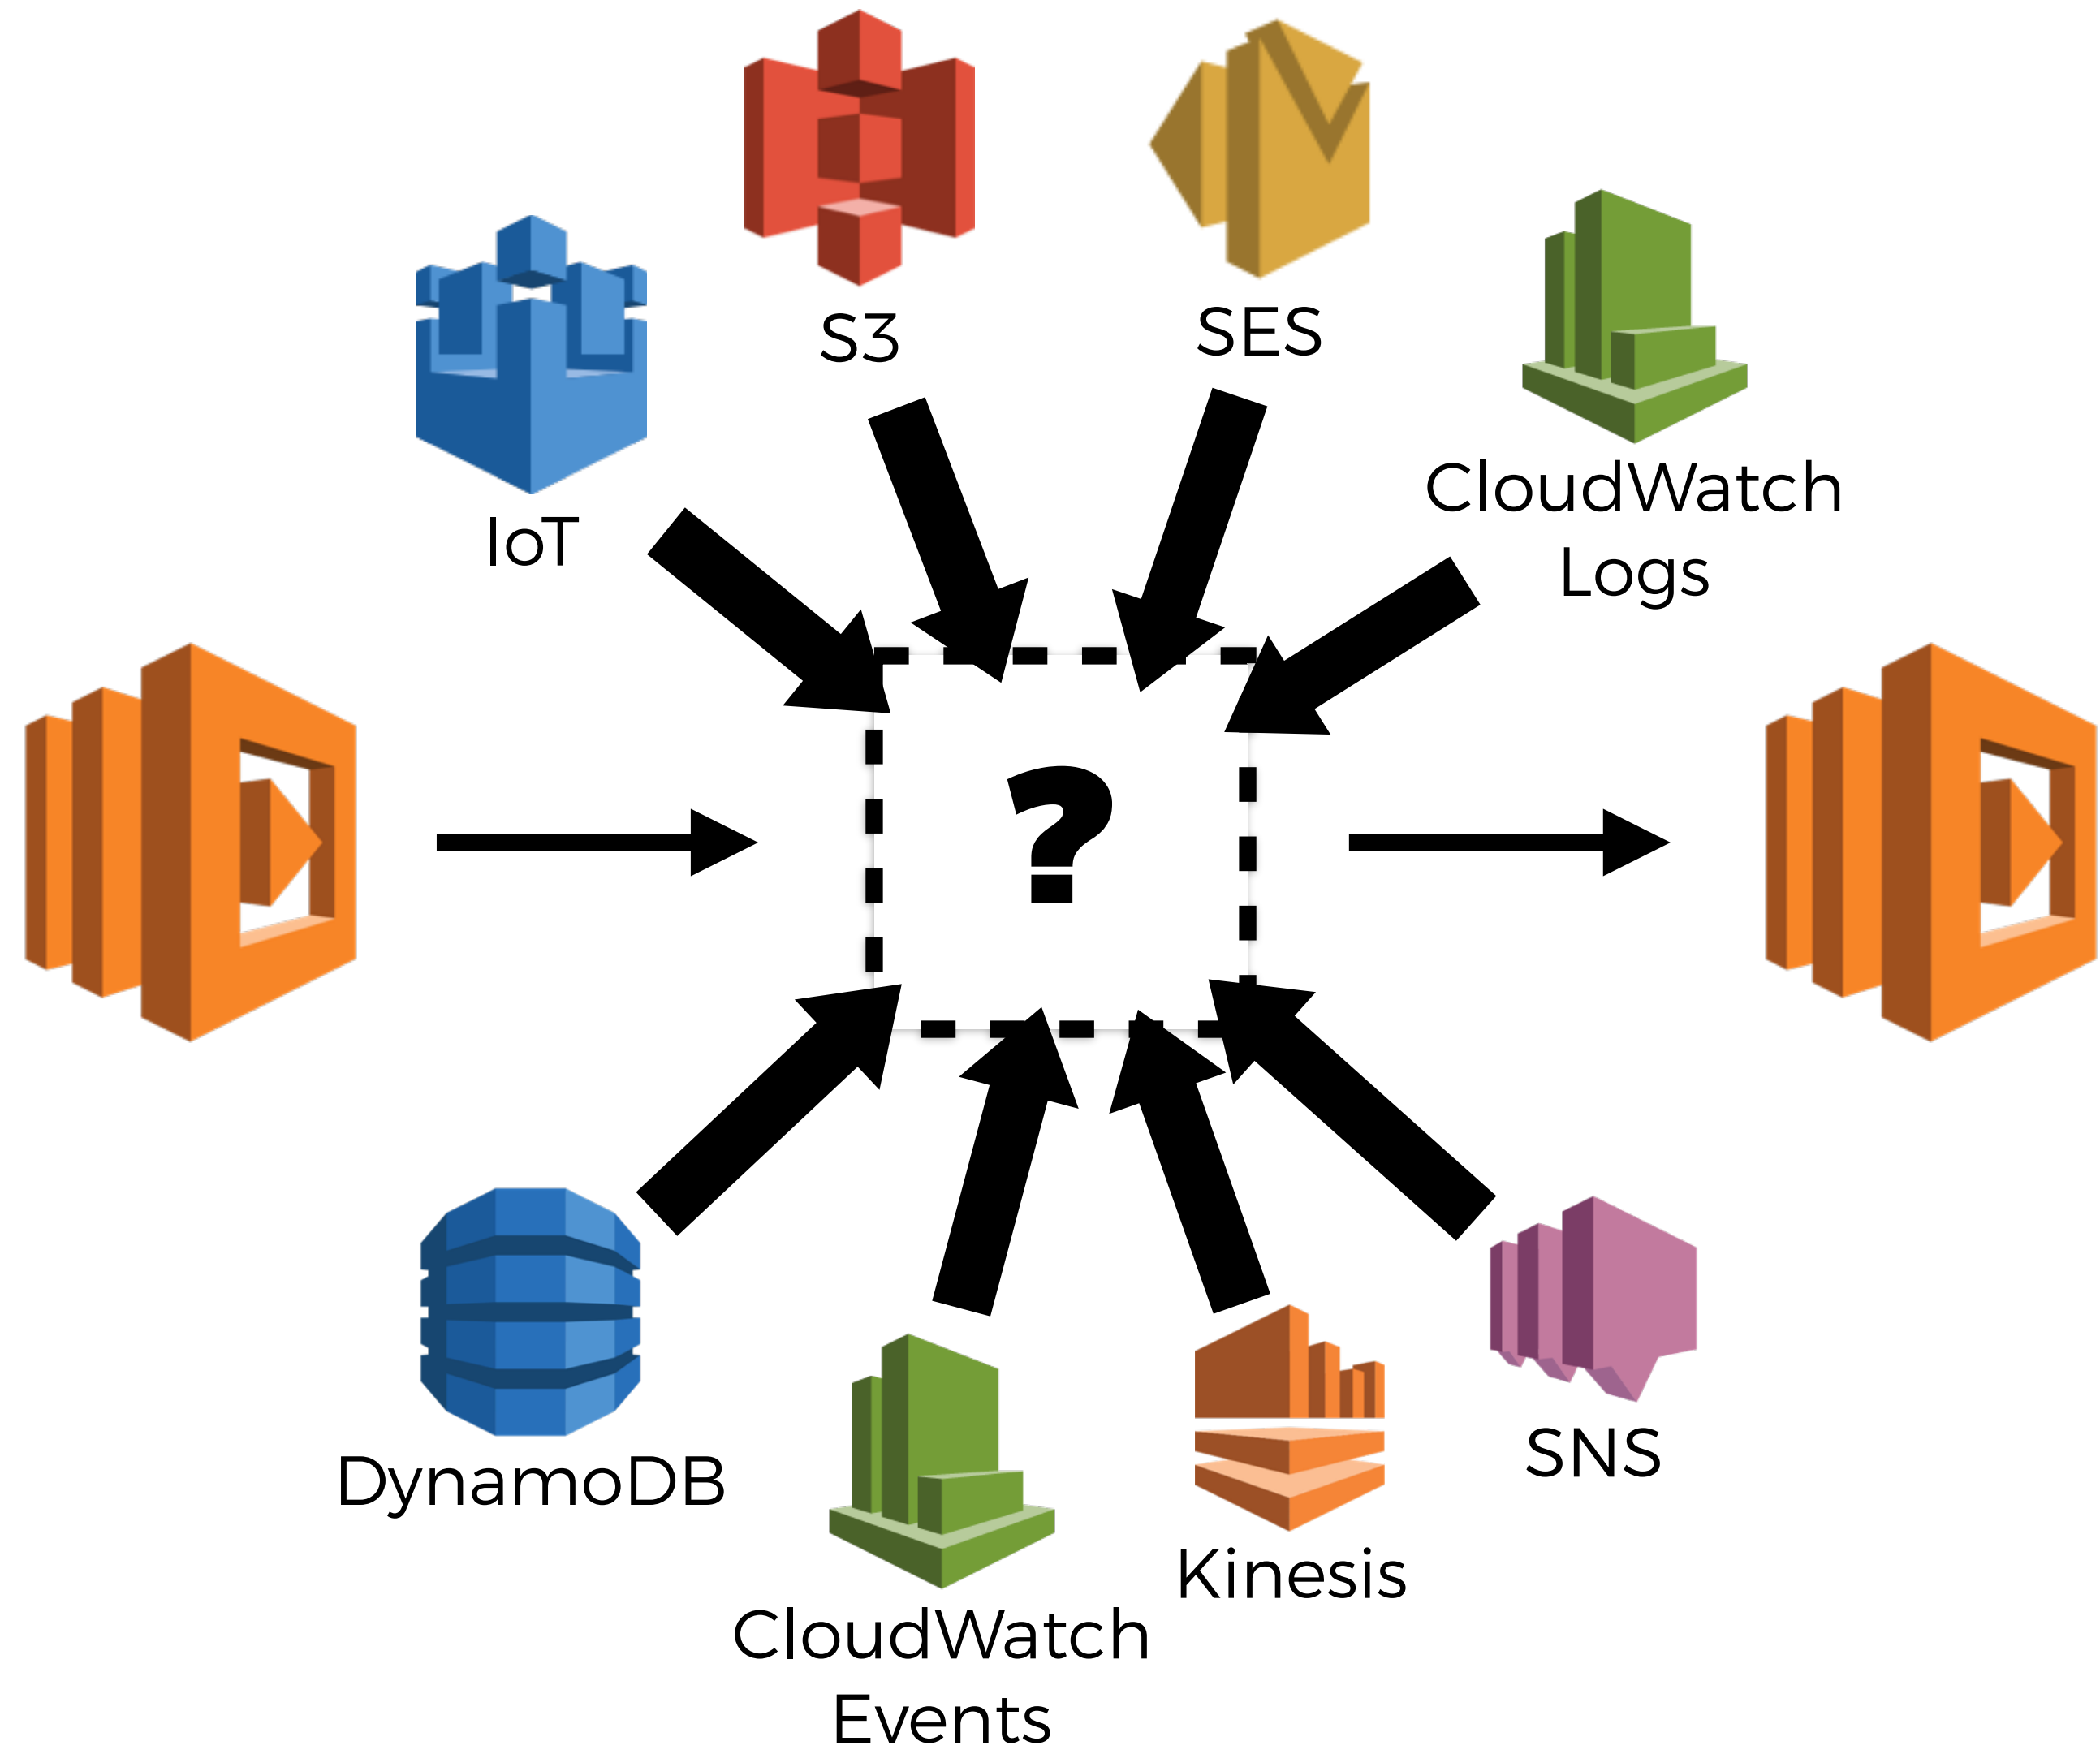
Events As First-Class Citizens

We've all heard of events and event-driven programming, but in my experience, events are not used nearly enough in our (software) lives. We often don't appreciate how powerful this tool can be in our toolbox, and consequently we don't take advantage of it when we really should.

<http://bit.ly/2Dpidje>

functions are often chained together
via **asynchronous** invocations







IoT



S3



SES



CloudWatch

tracing **ASYNCHRONOUS**
invocations through so many
different event sources is difficult



DynamoDB



CloudWatch
Events



Kinesis



SNS

new challenges


- **nowhere** to install agents/daemons
- **no** background processing
- **higher** concurrency to telemetry system
- **high** chance of data loss (if batching)
- **asynchronous** invocations



the Present



These are the four pillars of the Observability Engineering team's charter:

- *Monitoring*
- *Alerting/visualization*
- *Distributed systems tracing infrastructure*
- *Log aggregation/analytics* 

- Observability Engineering at Twitter <http://bit.ly/2DnjyuW>

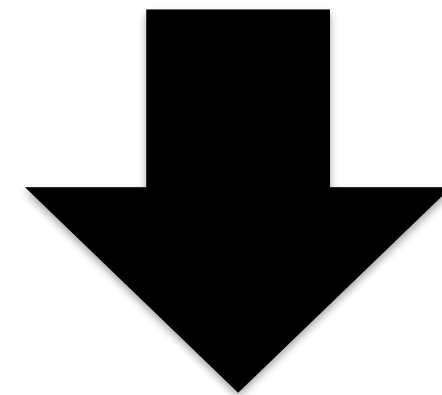




```
console.log("GOT is off air, what do I do now?");
```



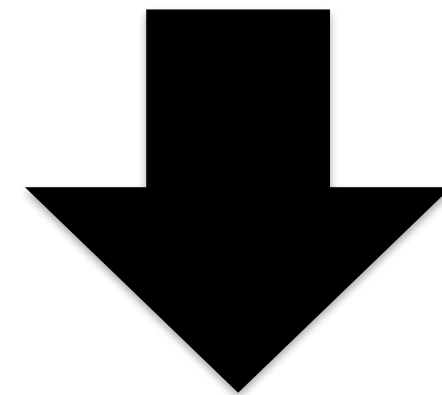

```
console.log("GOT is off air, what do I do now?");
```



```
2016-07-12T12:24:37.571Z 994f18f9-482b-11e6-8668-53e4eab441ae  
GOT is off air, what do I do now?
```



```
console.log("GOT is off air, what do I do now?");
```



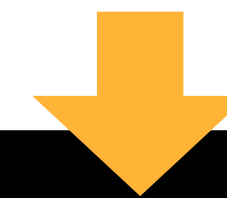
UTC Timestamp

Request Id



```
2016-07-12T12:24:37.571Z 994f18f9-482b-11e6-8668-53e4eab441ae  
GOT is off air, what do I do now?
```

your log message





Filter:

<input type="checkbox"/>	Log Streams	▼	Last Event Time	▼
<input type="checkbox"/>	2018/01/13/[\$LATEST]cd9779a0b20f44db8a661ca156dff8a1		2018-01-13 14:23 UTC	
<input type="checkbox"/>	2018/01/13/[\$LATEST]cb65dd255c9a426a8516ae38b06cd0bd		2018-01-13 13:55 UTC	
<input type="checkbox"/>	2018/01/13/[\$LATEST]064fbf40996847fd8a24b5b6cbd1e5af		2018-01-13 13:34 UTC	
<input type="checkbox"/>	2018/01/07/[\$LATEST]a60a31cdc718407bb7dd42b05f4735dc		2018-01-07 02:52 UTC	
<input type="checkbox"/>	2018/01/07/[\$LATEST]1c95f57d79624ef388349e71ce4ff34f		2018-01-07 02:48 UTC	
<input type="checkbox"/>	2018/01/07/[\$LATEST]bd9a05445760478b95d78afc916c4bd6		2018-01-07 01:23 UTC	
<input type="checkbox"/>	2018/01/07/[\$LATEST]00775bac85a04cb58bb421b87f6d57f5		2018-01-07 01:17 UTC	
<input type="checkbox"/>	2018/01/06/[\$LATEST]5c55f992924643718bfec831df799d3f		2018-01-06 22:27 UTC	
<input type="checkbox"/>	2018/01/06/[\$LATEST]45162e45baf64db697de04bb4af96aa1		2018-01-06 13:39 UTC	
<input type="checkbox"/>	2018/01/06/[\$LATEST]63b8533fec2845e8b65af4505984afa5		2018-01-06 03:57 UTC	
<input type="checkbox"/>	2018/01/06/[\$LATEST]63ff65ac546440889e944074f2b3a1e9		2018-01-06 01:24 UTC	
<input type="checkbox"/>	2018/01/02/[\$LATEST]f264f6c6cf344b4c8d6e9f549bdfc26c		2018-01-02 19:38 UTC	
<input type="checkbox"/>	2018/01/02/[\$LATEST]07308effe6dd4f549d23bd8adc5a8f93		2018-01-02 15:20 UTC	
<input type="checkbox"/>	2017/12/24/[\$LATEST]c3e00904081b4477ab254d823734c28c		2017-12-24 03:39 UTC	
<input type="checkbox"/>	2017/12/24/[\$LATEST]bdfbfe1fbf9d45a7a77eae39742d4165		2017-12-24 03:38 UTC	
<input type="checkbox"/>	2017/12/24/[\$LATEST]24249247d9864f8b954eea0554bd625d		2017-12-24 03:35 UTC	
<input type="checkbox"/>	2017/12/24/[\$LATEST]63851da58bcb4b078ceee2098e287c96		2017-12-24 03:31 UTC	
<input type="checkbox"/>	2017/12/24/[\$LATEST]16ca91710e521d11e20b90b86319573e		2017-12-24 03:24 UTC	



Search Log Group Create Log Stream Delete Log Stream

Filter: Log Stream Name Prefix

one **log group** per function

<input type="checkbox"/>	Log Streams	Last Event Time
<input type="checkbox"/>	2018/01/13/[\$LATEST]cd9779a0b20f44db8a661ca156dff8a1	2018-01-13 14:23 UTC
<input type="checkbox"/>	2018/01/13/[\$LATEST]cb65dd255c9a426a8516ae38b06cd0bd	2018-01-13 13:55 UTC
<input type="checkbox"/>	2018/01/13/[\$LATEST]064fbf40996847fd8a24b5b6cbd1e5af	2018-01-13 13:34 UTC
<input type="checkbox"/>	2018/01/07/[\$LATEST]a60a31cdc718407bb7dd42b05f4735dc	2018-01-07 02:52 UTC
<input type="checkbox"/>	2018/01/07/[\$LATEST]1c95f57d79624ef388349e71ce4ff34f	
<input type="checkbox"/>	2018/01/07/[\$LATEST]bd9a05445760478b95d78afc916c4bd6	
<input type="checkbox"/>	2018/01/07/[\$LATEST]00775bac85a04cb58bb421b87f6d57f5	
<input type="checkbox"/>	2018/01/06/[\$LATEST]5c55f992924643718bfec831df799d3f	
<input type="checkbox"/>	2018/01/06/[\$LATEST]45162e45baf64db697de04bb4af96aa1	
<input type="checkbox"/>	2018/01/06/[\$LATEST]63b8533fec2845e8b65af4505984afa5	2018-01-06 03:57 UTC
<input type="checkbox"/>	2018/01/06/[\$LATEST]63ff65ac546440889e944074f2b3a1e9	2018-01-06 01:24 UTC
<input type="checkbox"/>	2018/01/02/[\$LATEST]f264f6c6cf344b4c8d6e9f549bdfc26c	2018-01-02 19:38 UTC
<input type="checkbox"/>	2018/01/02/[\$LATEST]07308effe6dd4f549d23bd8adc5a8f93	2018-01-02 15:20 UTC
<input type="checkbox"/>	2017/12/24/[\$LATEST]c3e00904081b4477ab254d823734c28c	2017-12-24 03:39 UTC
<input type="checkbox"/>	2017/12/24/[\$LATEST]bdfbfe1fbf9d45a7a77eae39742d4165	2017-12-24 03:38 UTC
<input type="checkbox"/>	2017/12/24/[\$LATEST]24249247d9864f8b954eea0554bd625d	2017-12-24 03:35 UTC
<input type="checkbox"/>	2017/12/24/[\$LATEST]63851da58bcb4b078ceee2098e287c96	2017-12-24 03:31 UTC
<input type="checkbox"/>	2017/12/24/[\$LATEST]16ca91710e521d11e20b90b86319573e	2017-12-24 03:24 UTC

one **log stream** for each concurrent invocation

logs are not easily searchable in
CloudWatch Logs



me



CloudWatch Logs

Create Metric Filter

Actions ▾

Filter: /aws/lambda/E

Log Groups

• /aws/lambda/ELBL

Create log group

Delete log group

Export

Export data to Amazon S3

View all exports to Amazon S3

Subscriptions

Stream to AWS Lambda

Stream to Amazon Elasticsearch Service

Remove Subscription Filter



CloudWatch Logs



AWS Lambda



ELK stack

Event Source

Build or customize an Event Pattern or set a Schedule to invoke Targets.

Event Pattern ⓘ Schedule ⓘ

Build event pattern to match events by service

Service Name

Event Type

Any operation Specific operation(s)

Event Pattern Preview

[Copy to clipboard](#) [Edit](#)

```
{
  "source": [
    "aws.logs"
  ],
  "detail-type": [
    "AWS API Call via CloudTrail"
  ],
  "detail": {
    "eventSource": [
      "logs.amazonaws.com"
    ],
    "eventName": [
      "CreateLogGroup"
    ]
  }
}
```

Targets

Select Target to invoke when an event matches your Event Pattern or when schedule is triggered.

Lambda function

Function

▶ Configure version/alias

▶ Configure input

[+ Add target*](#)

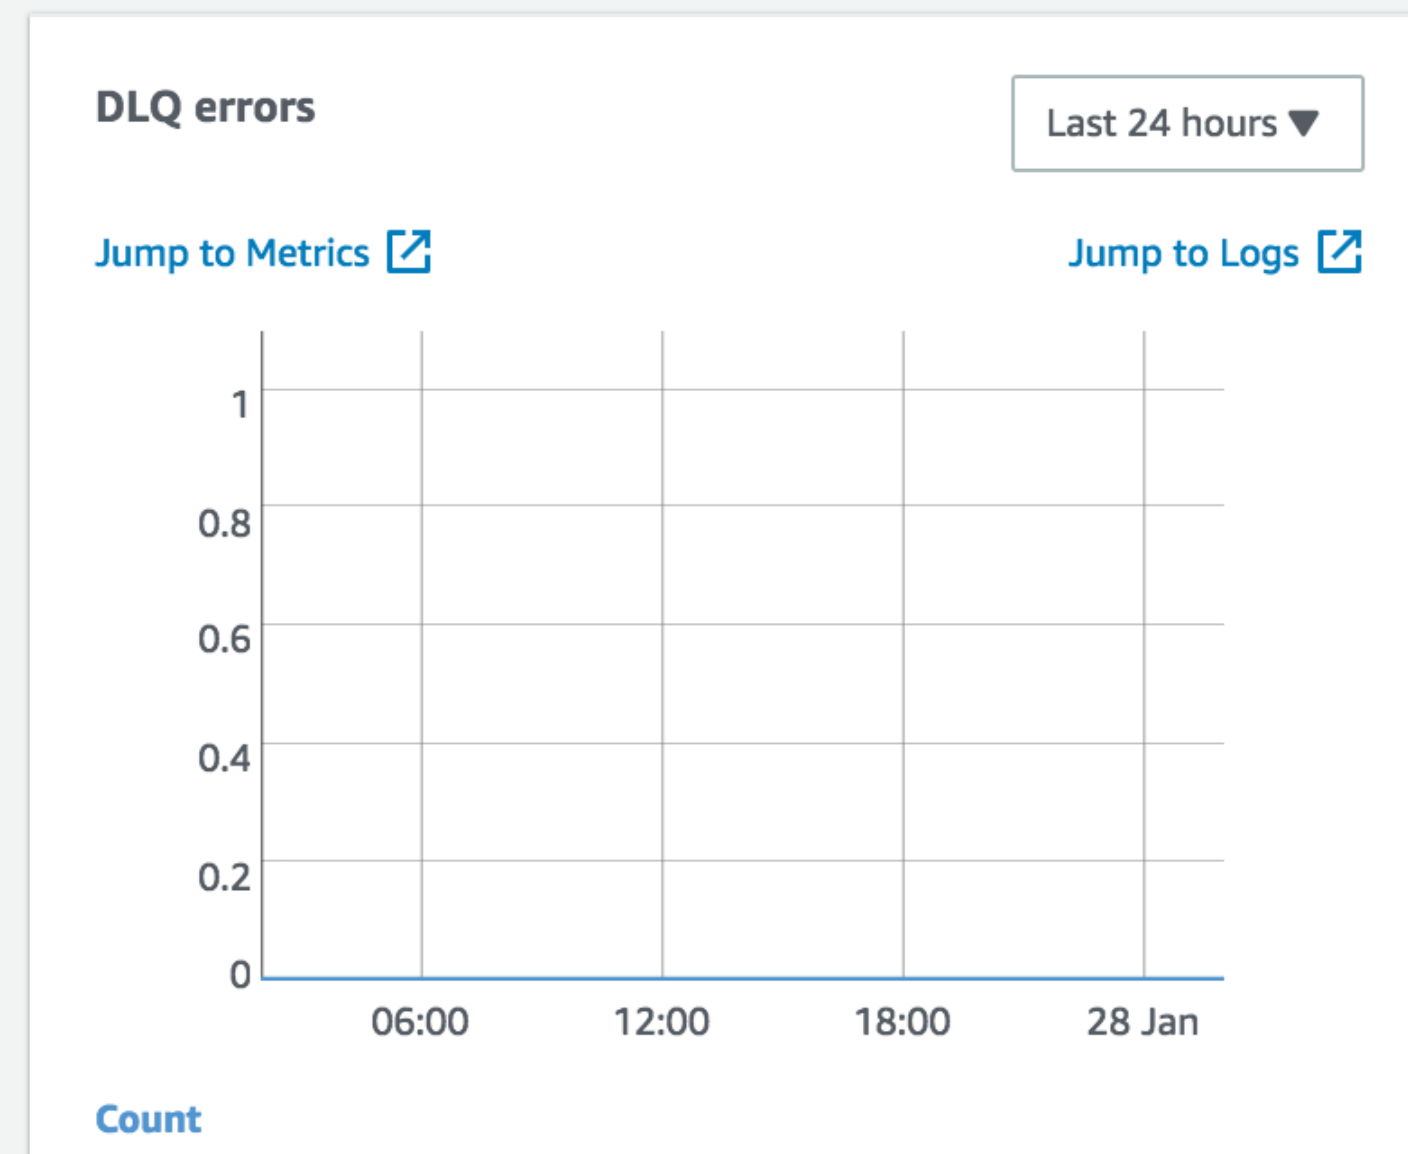
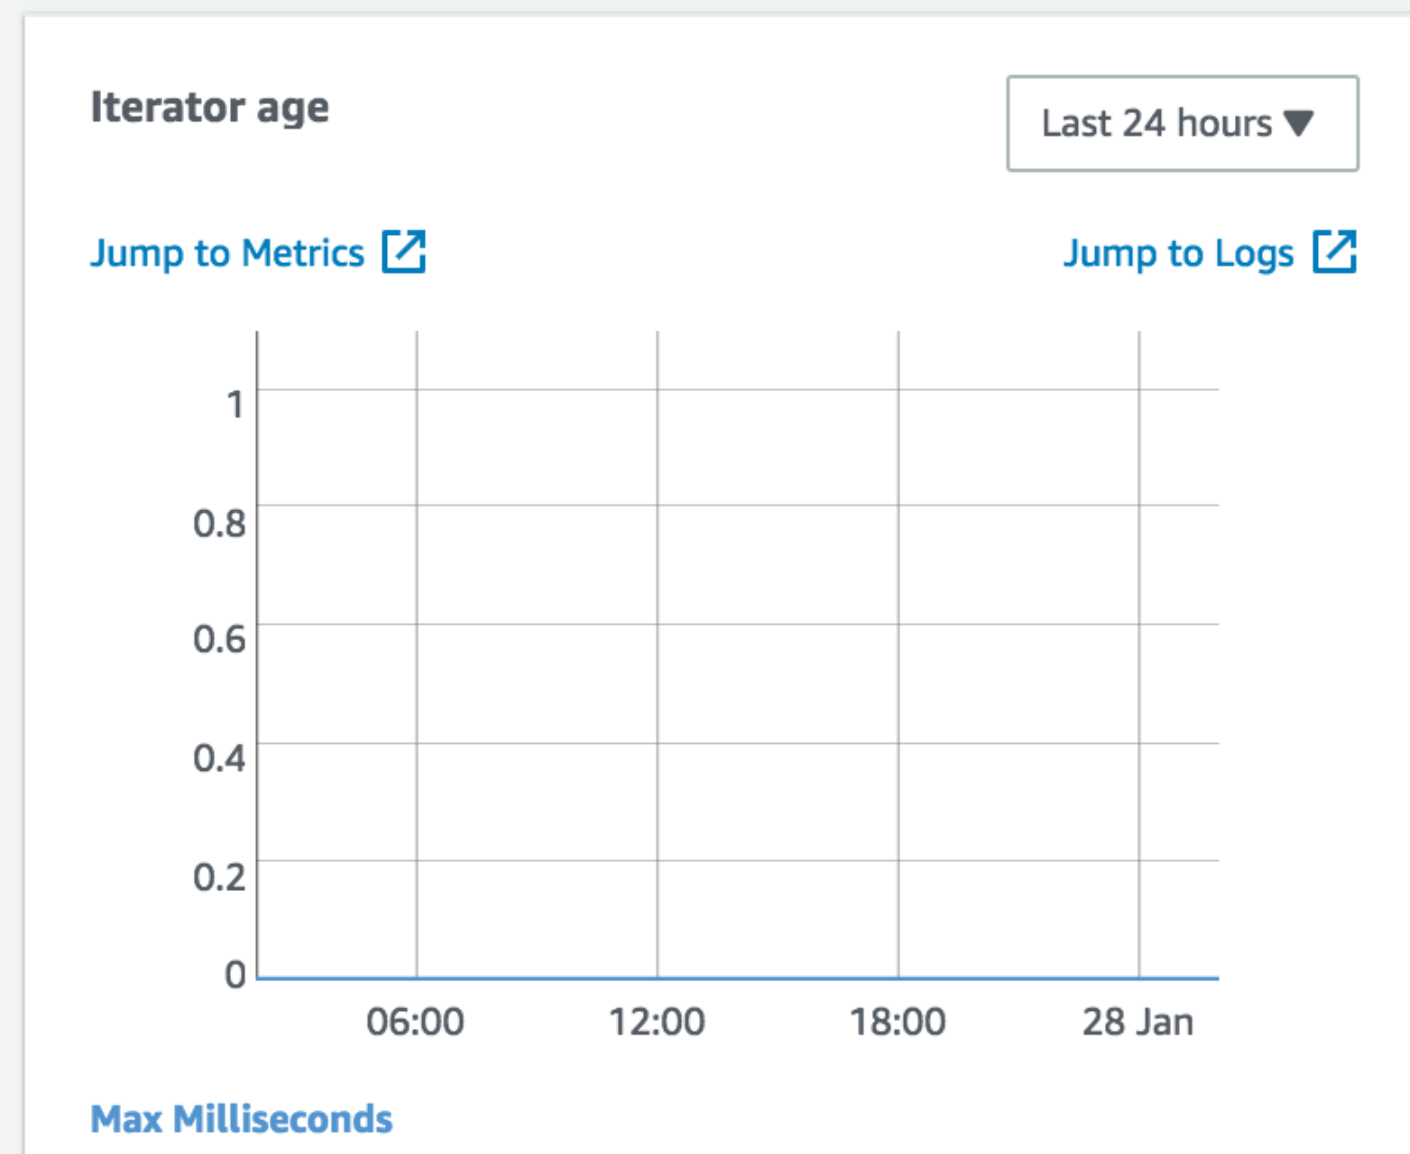
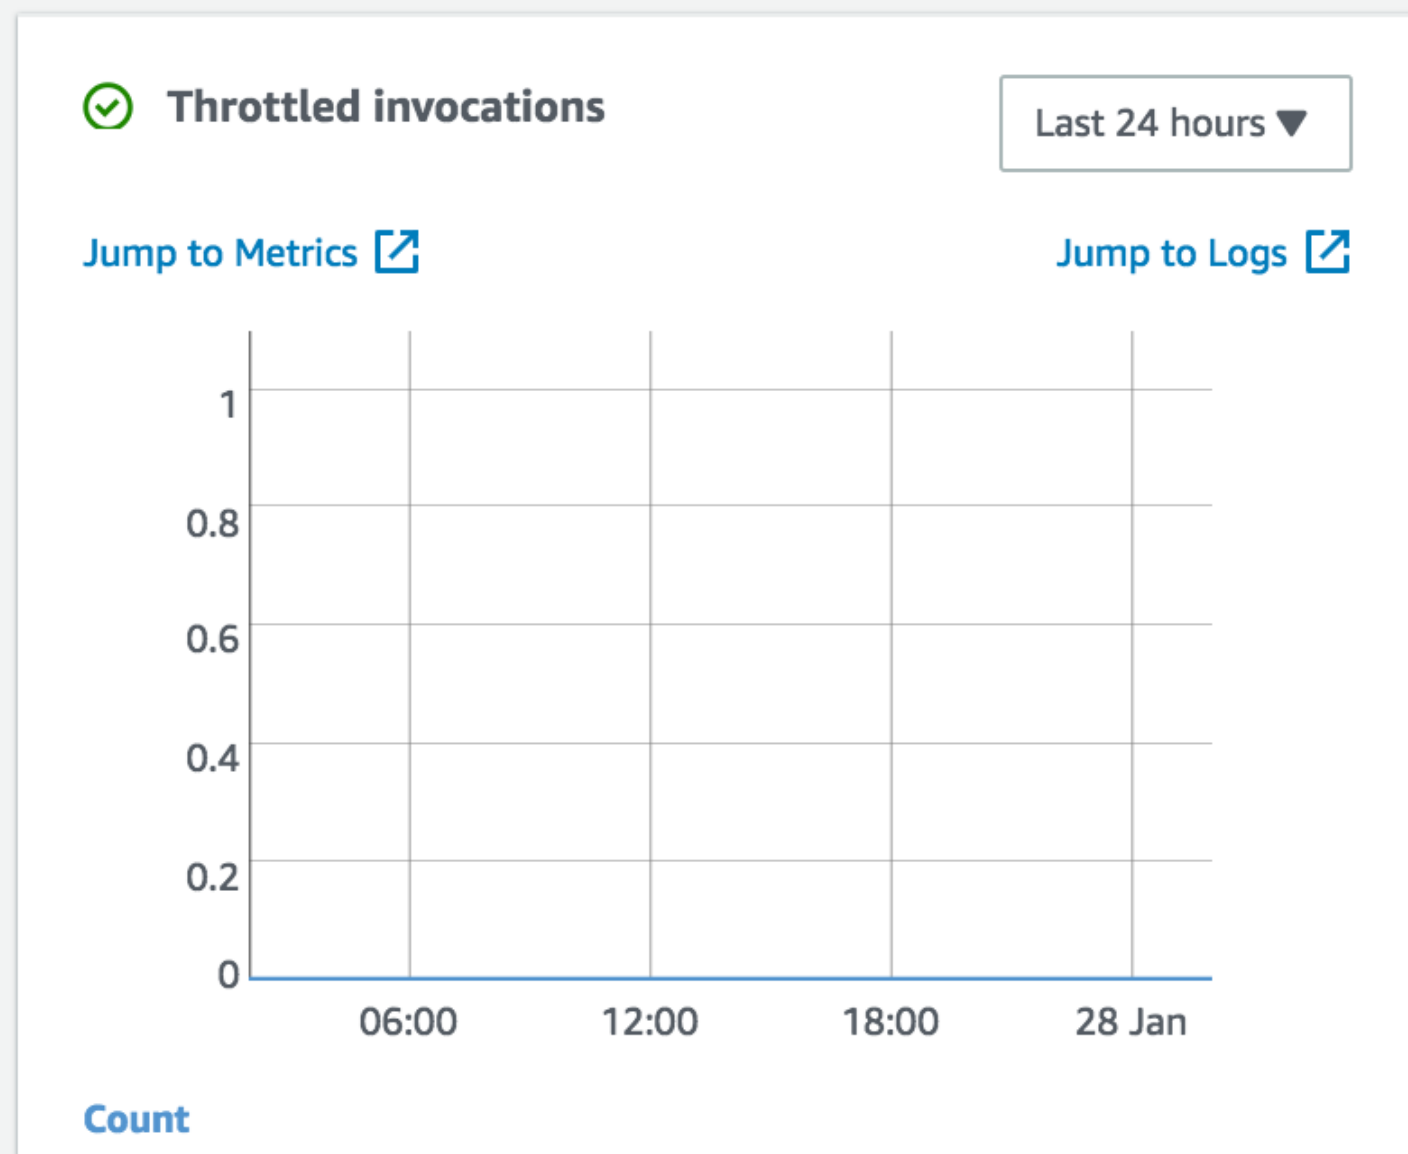
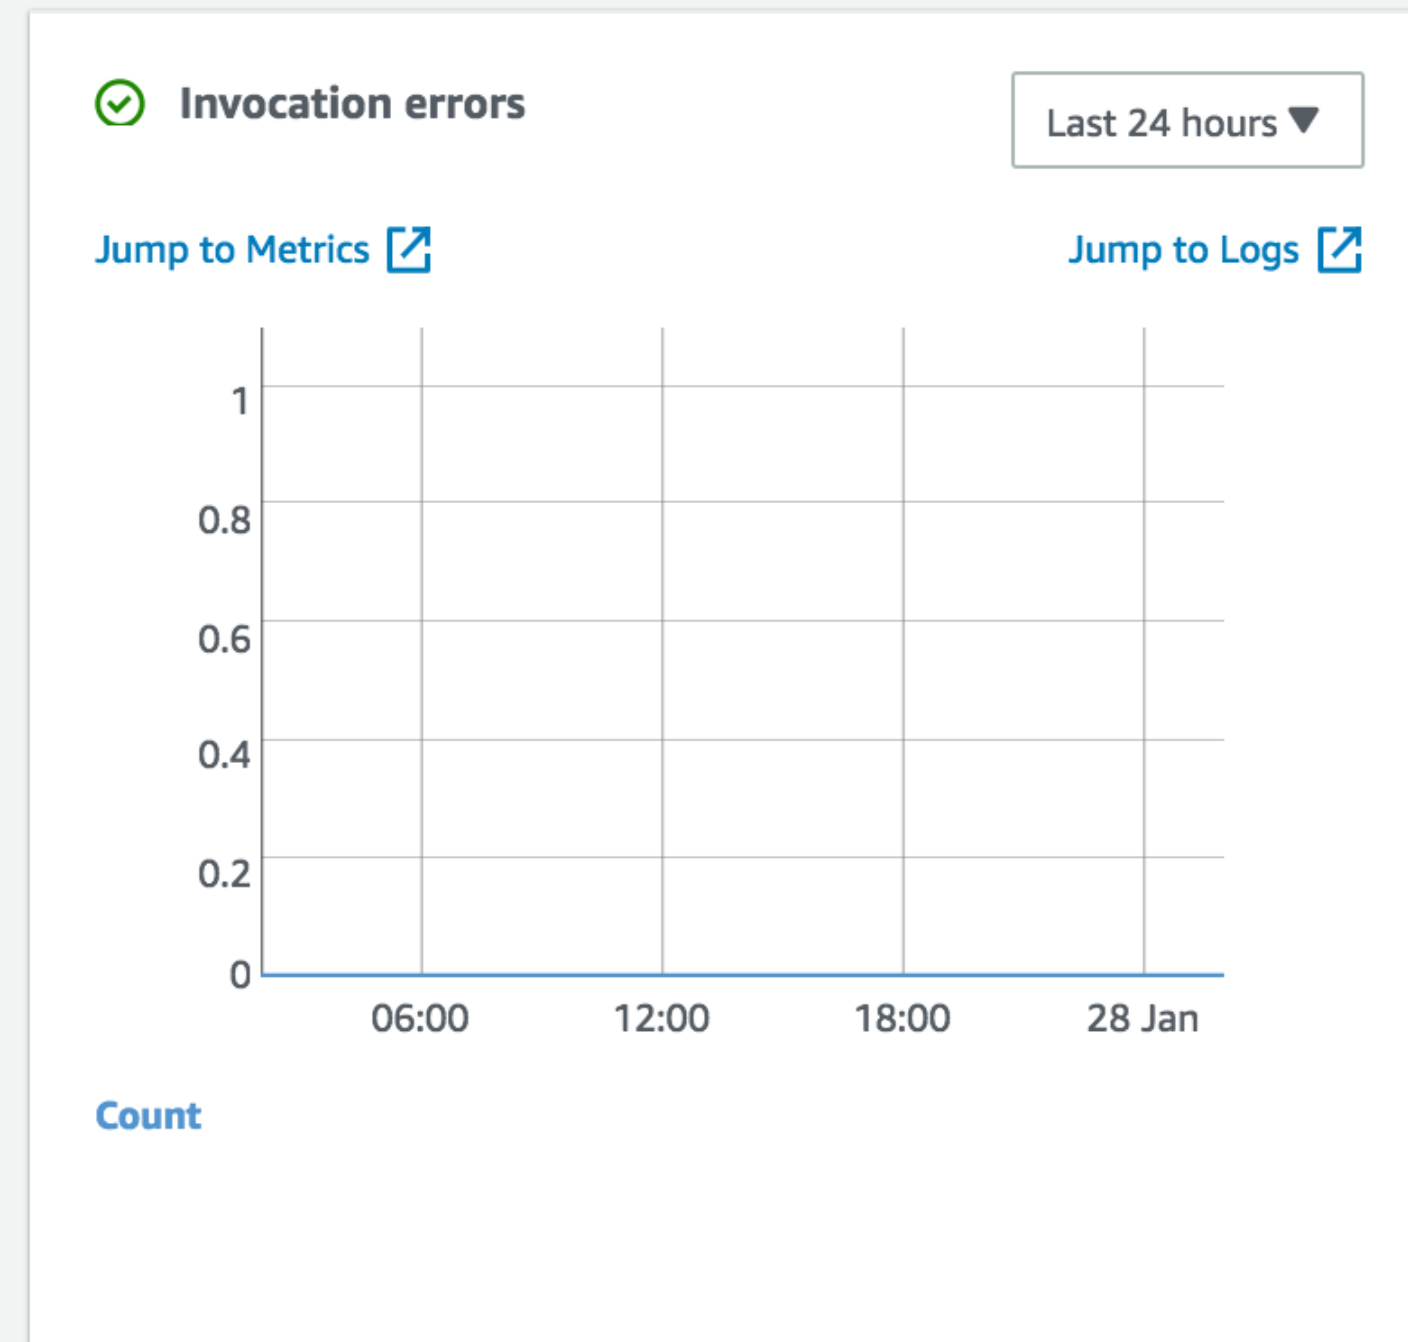


CloudWatch Logs

```
subscribe:
  handler: functions/subscribe/handler.handler
  description: Subscribe logs to the ship-log function
  environment:
    DEST_FUNC: "arn:aws:lambda:#{AWS::Region}:#{AWS::AccountId}:function:${self:service}-${se
  events:
    - cloudwatchEvent:
      event:
        source:
          - aws.logs
        detail-type:
          - AWS API Call via CloudTrail
        detail:
          eventSource:
            - logs.amazonaws.com
          eventName:
            - CreateLogGroup
```

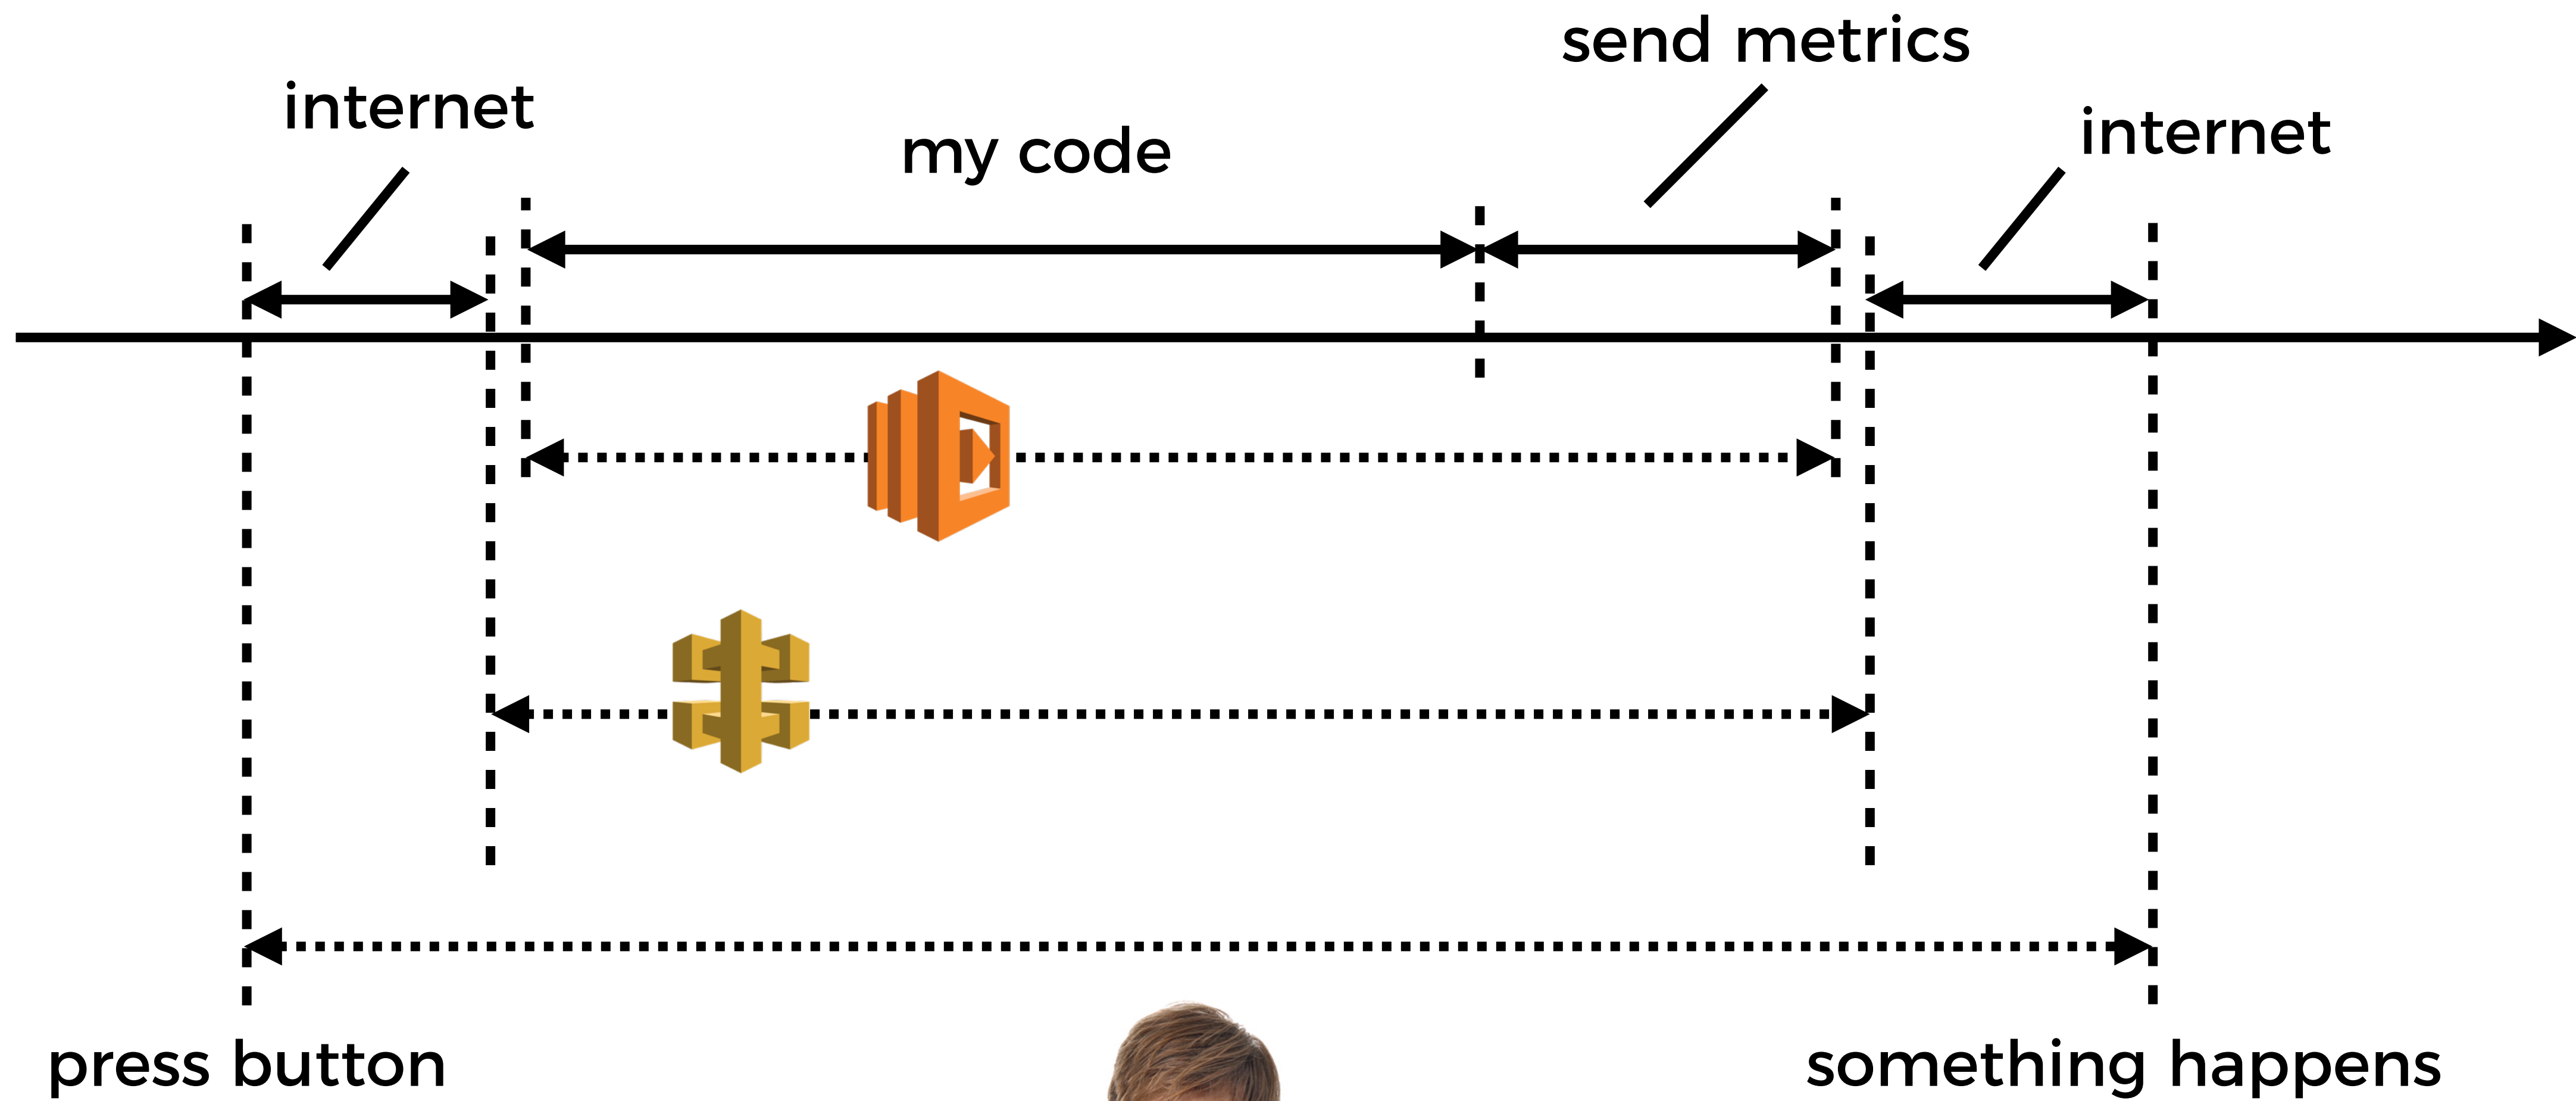

apig-iscoldstart-dev-hello

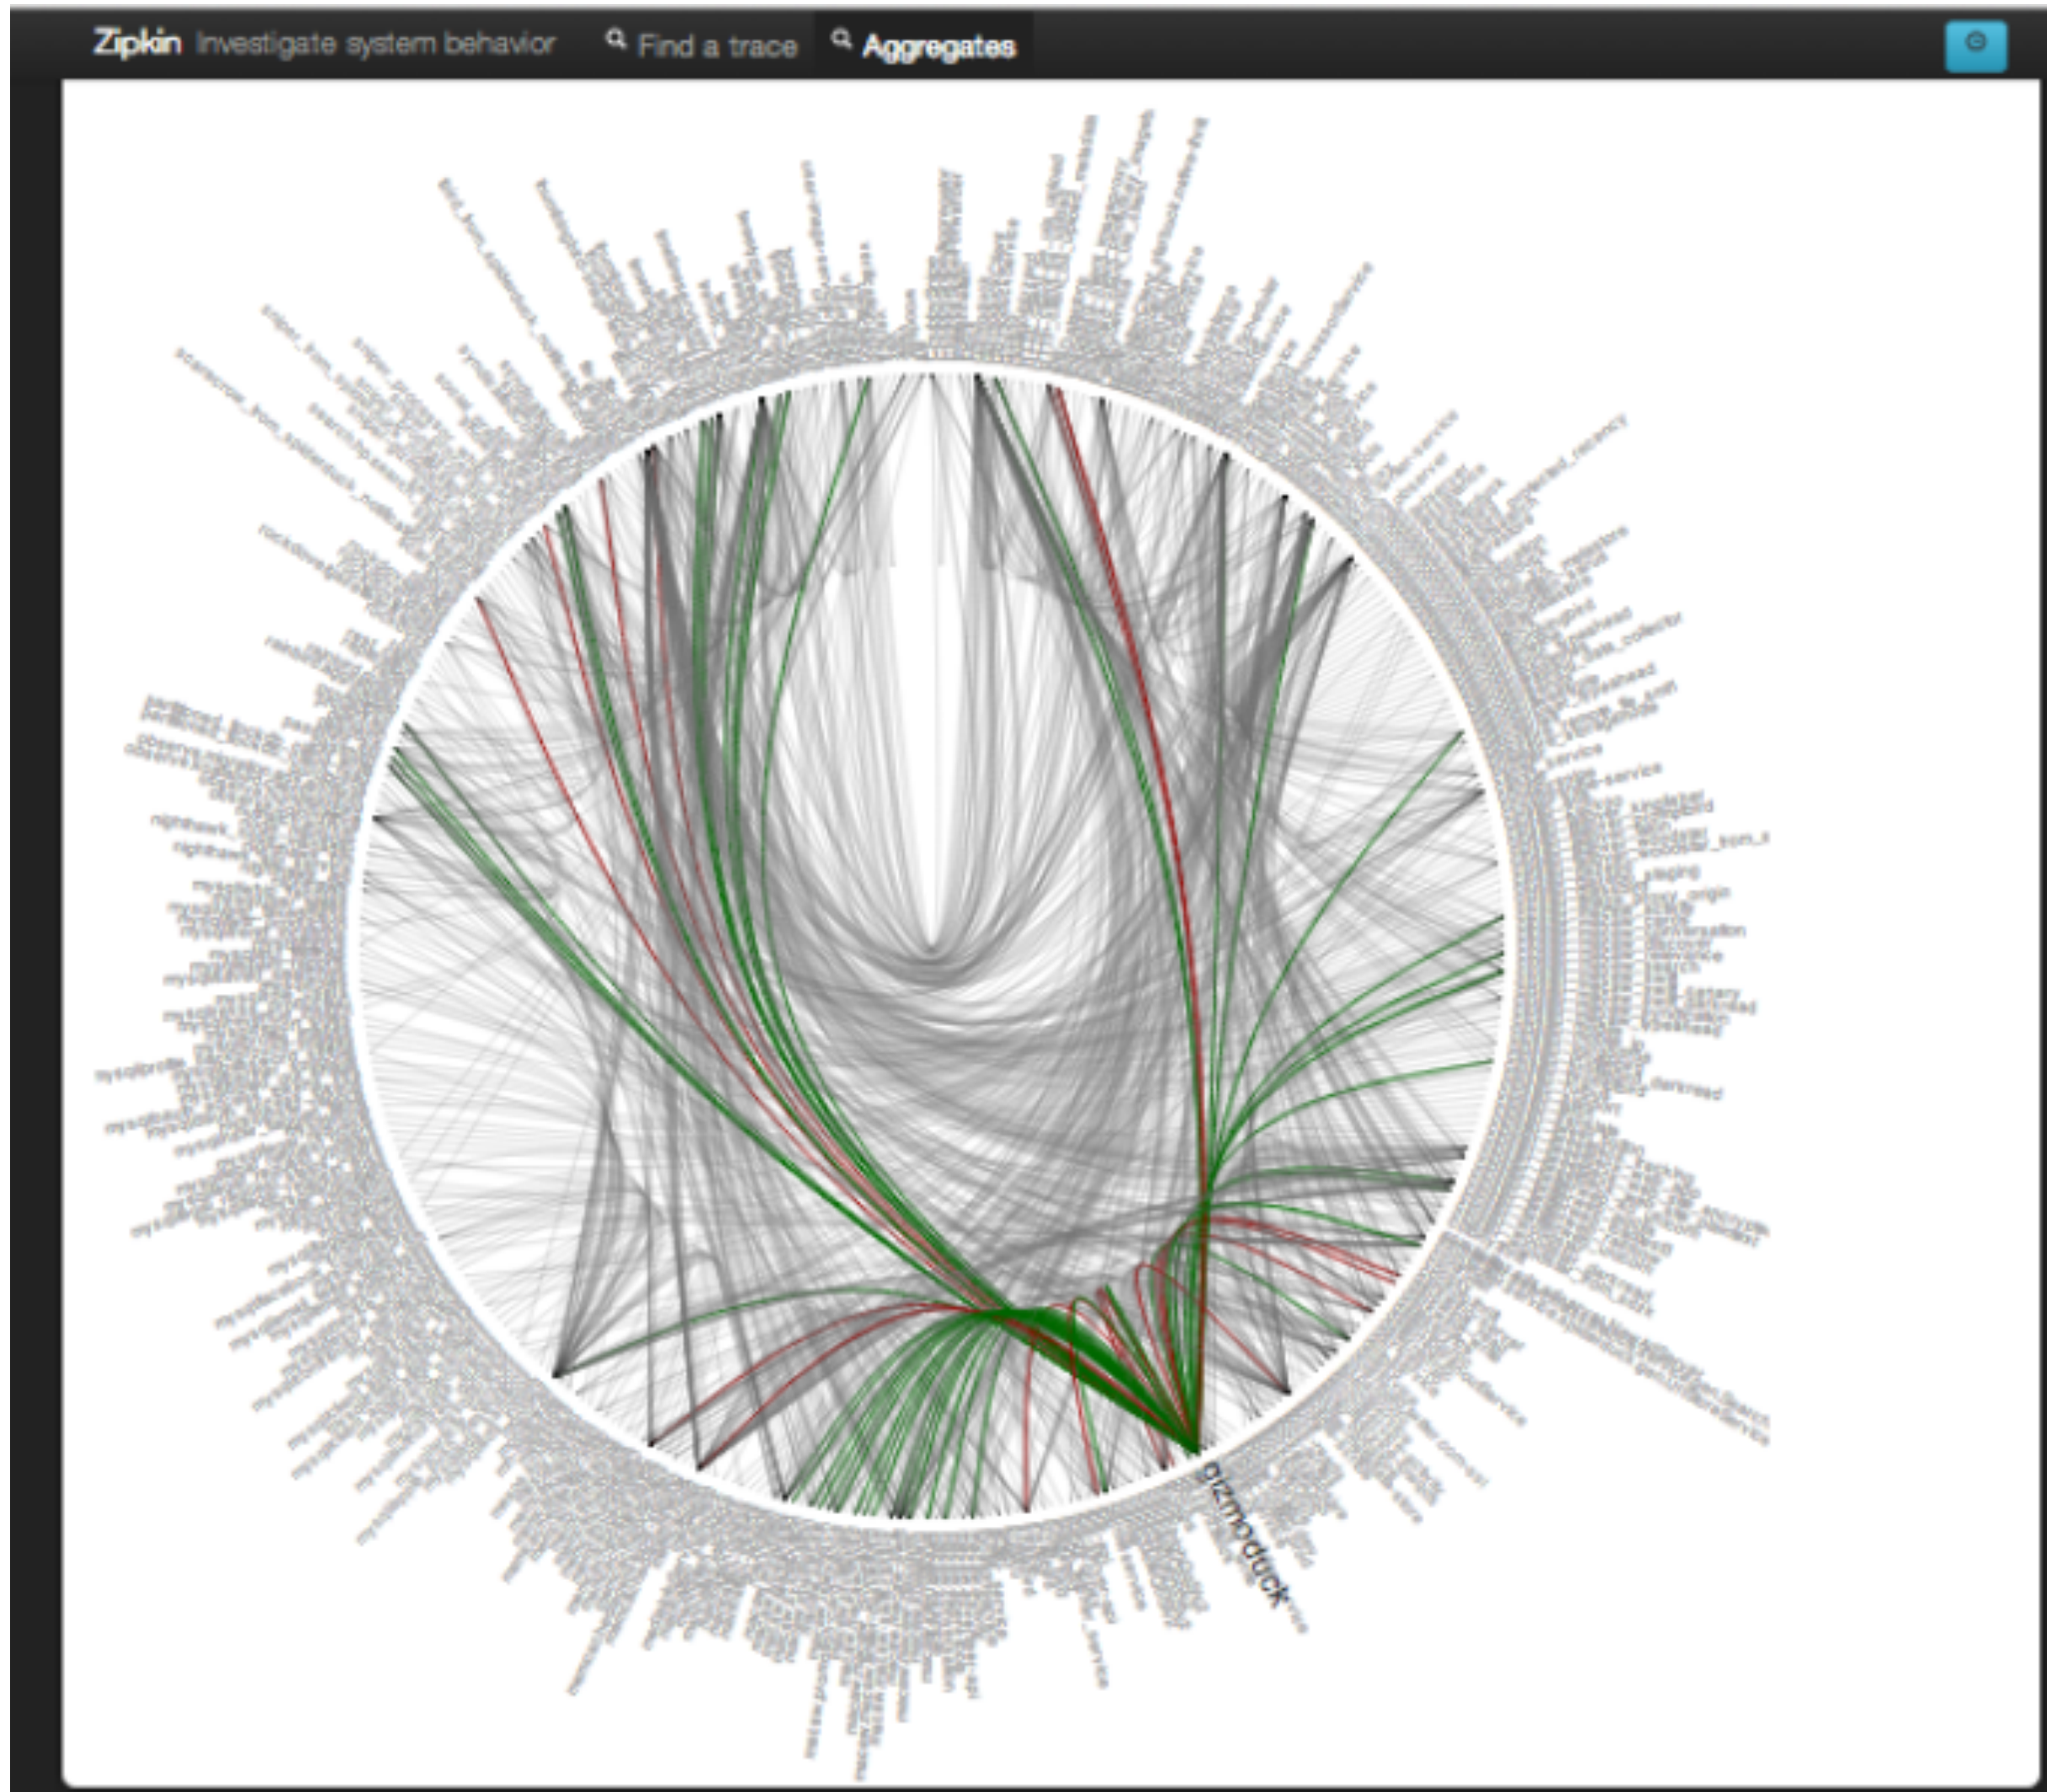
Qualifiers Actions Select a test event.. Test Save



new challenges

- **nowhere** to install agents/daemons
- **no** background processing





those extra 10-20ms for sending custom metrics would **compound** when you have microservices and multiple APIs are called within one slice of user event

Amazon found every **100ms** of latency **cost them 1% in sales.**

<http://bit.ly/2EXPfbA>

metrics

timestamp

metric value

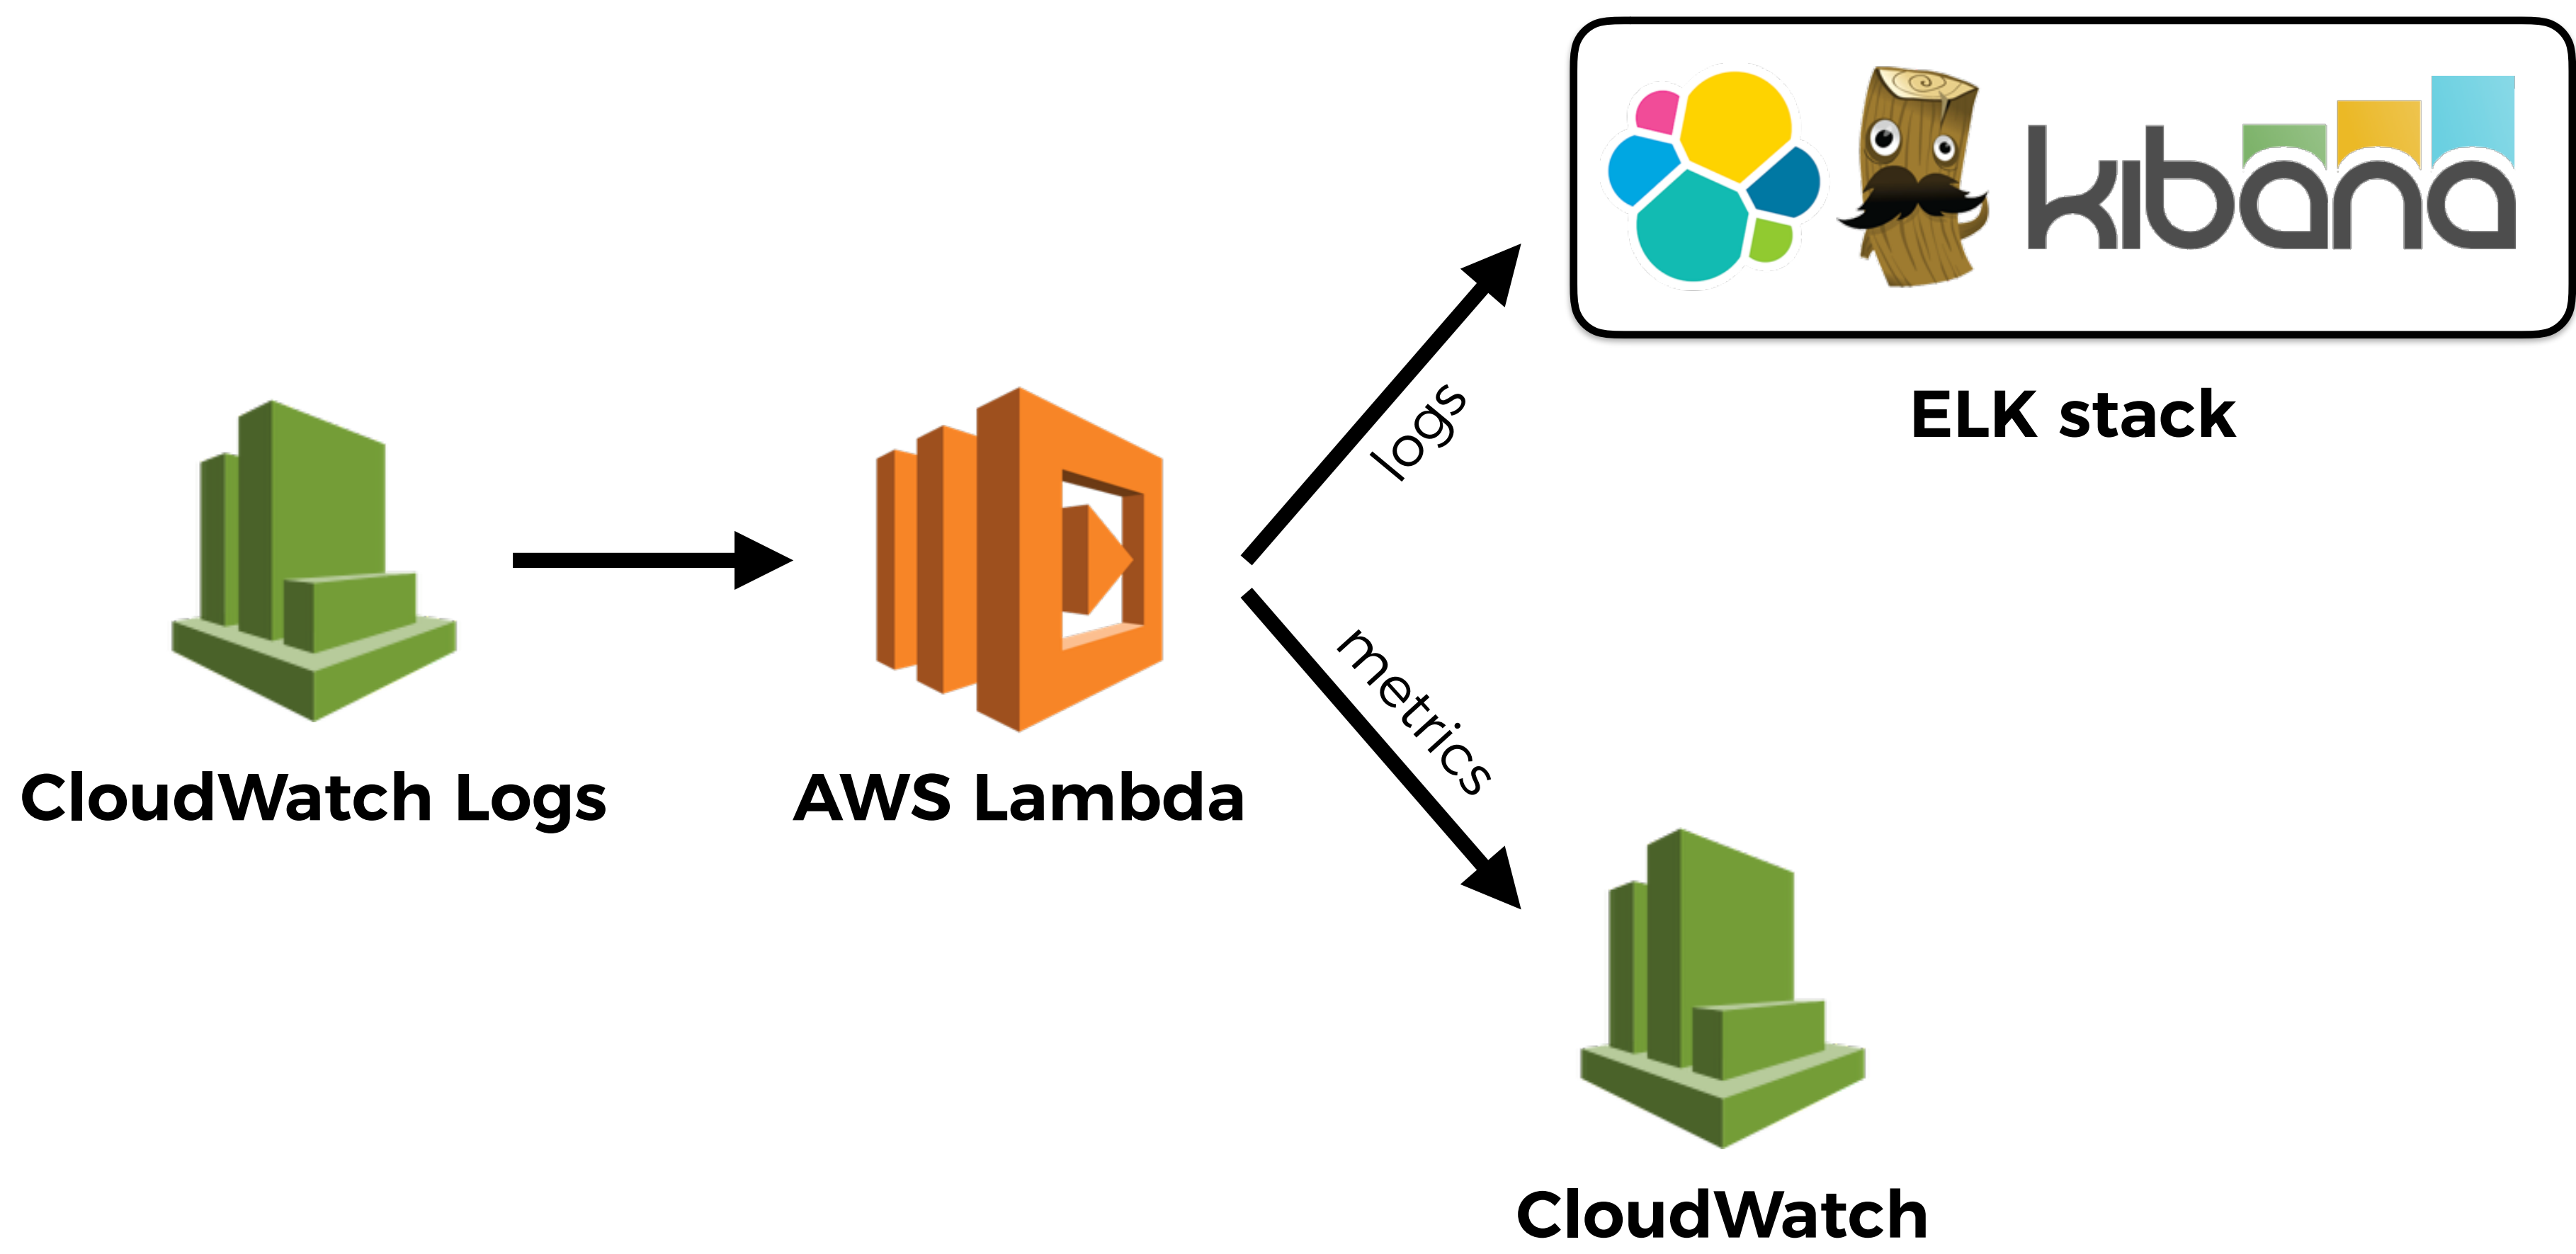
metric name

```
console.log("MONITORING | 1489795335 | 27.4 | latency | user-api-latency");  
console.log("MONITORING | 1489795335 | 8 | count | yubls-served");
```

logs

metric type

```
console.log("hydrating yubls from db...");  
console.log("fetching user info from user-api");
```

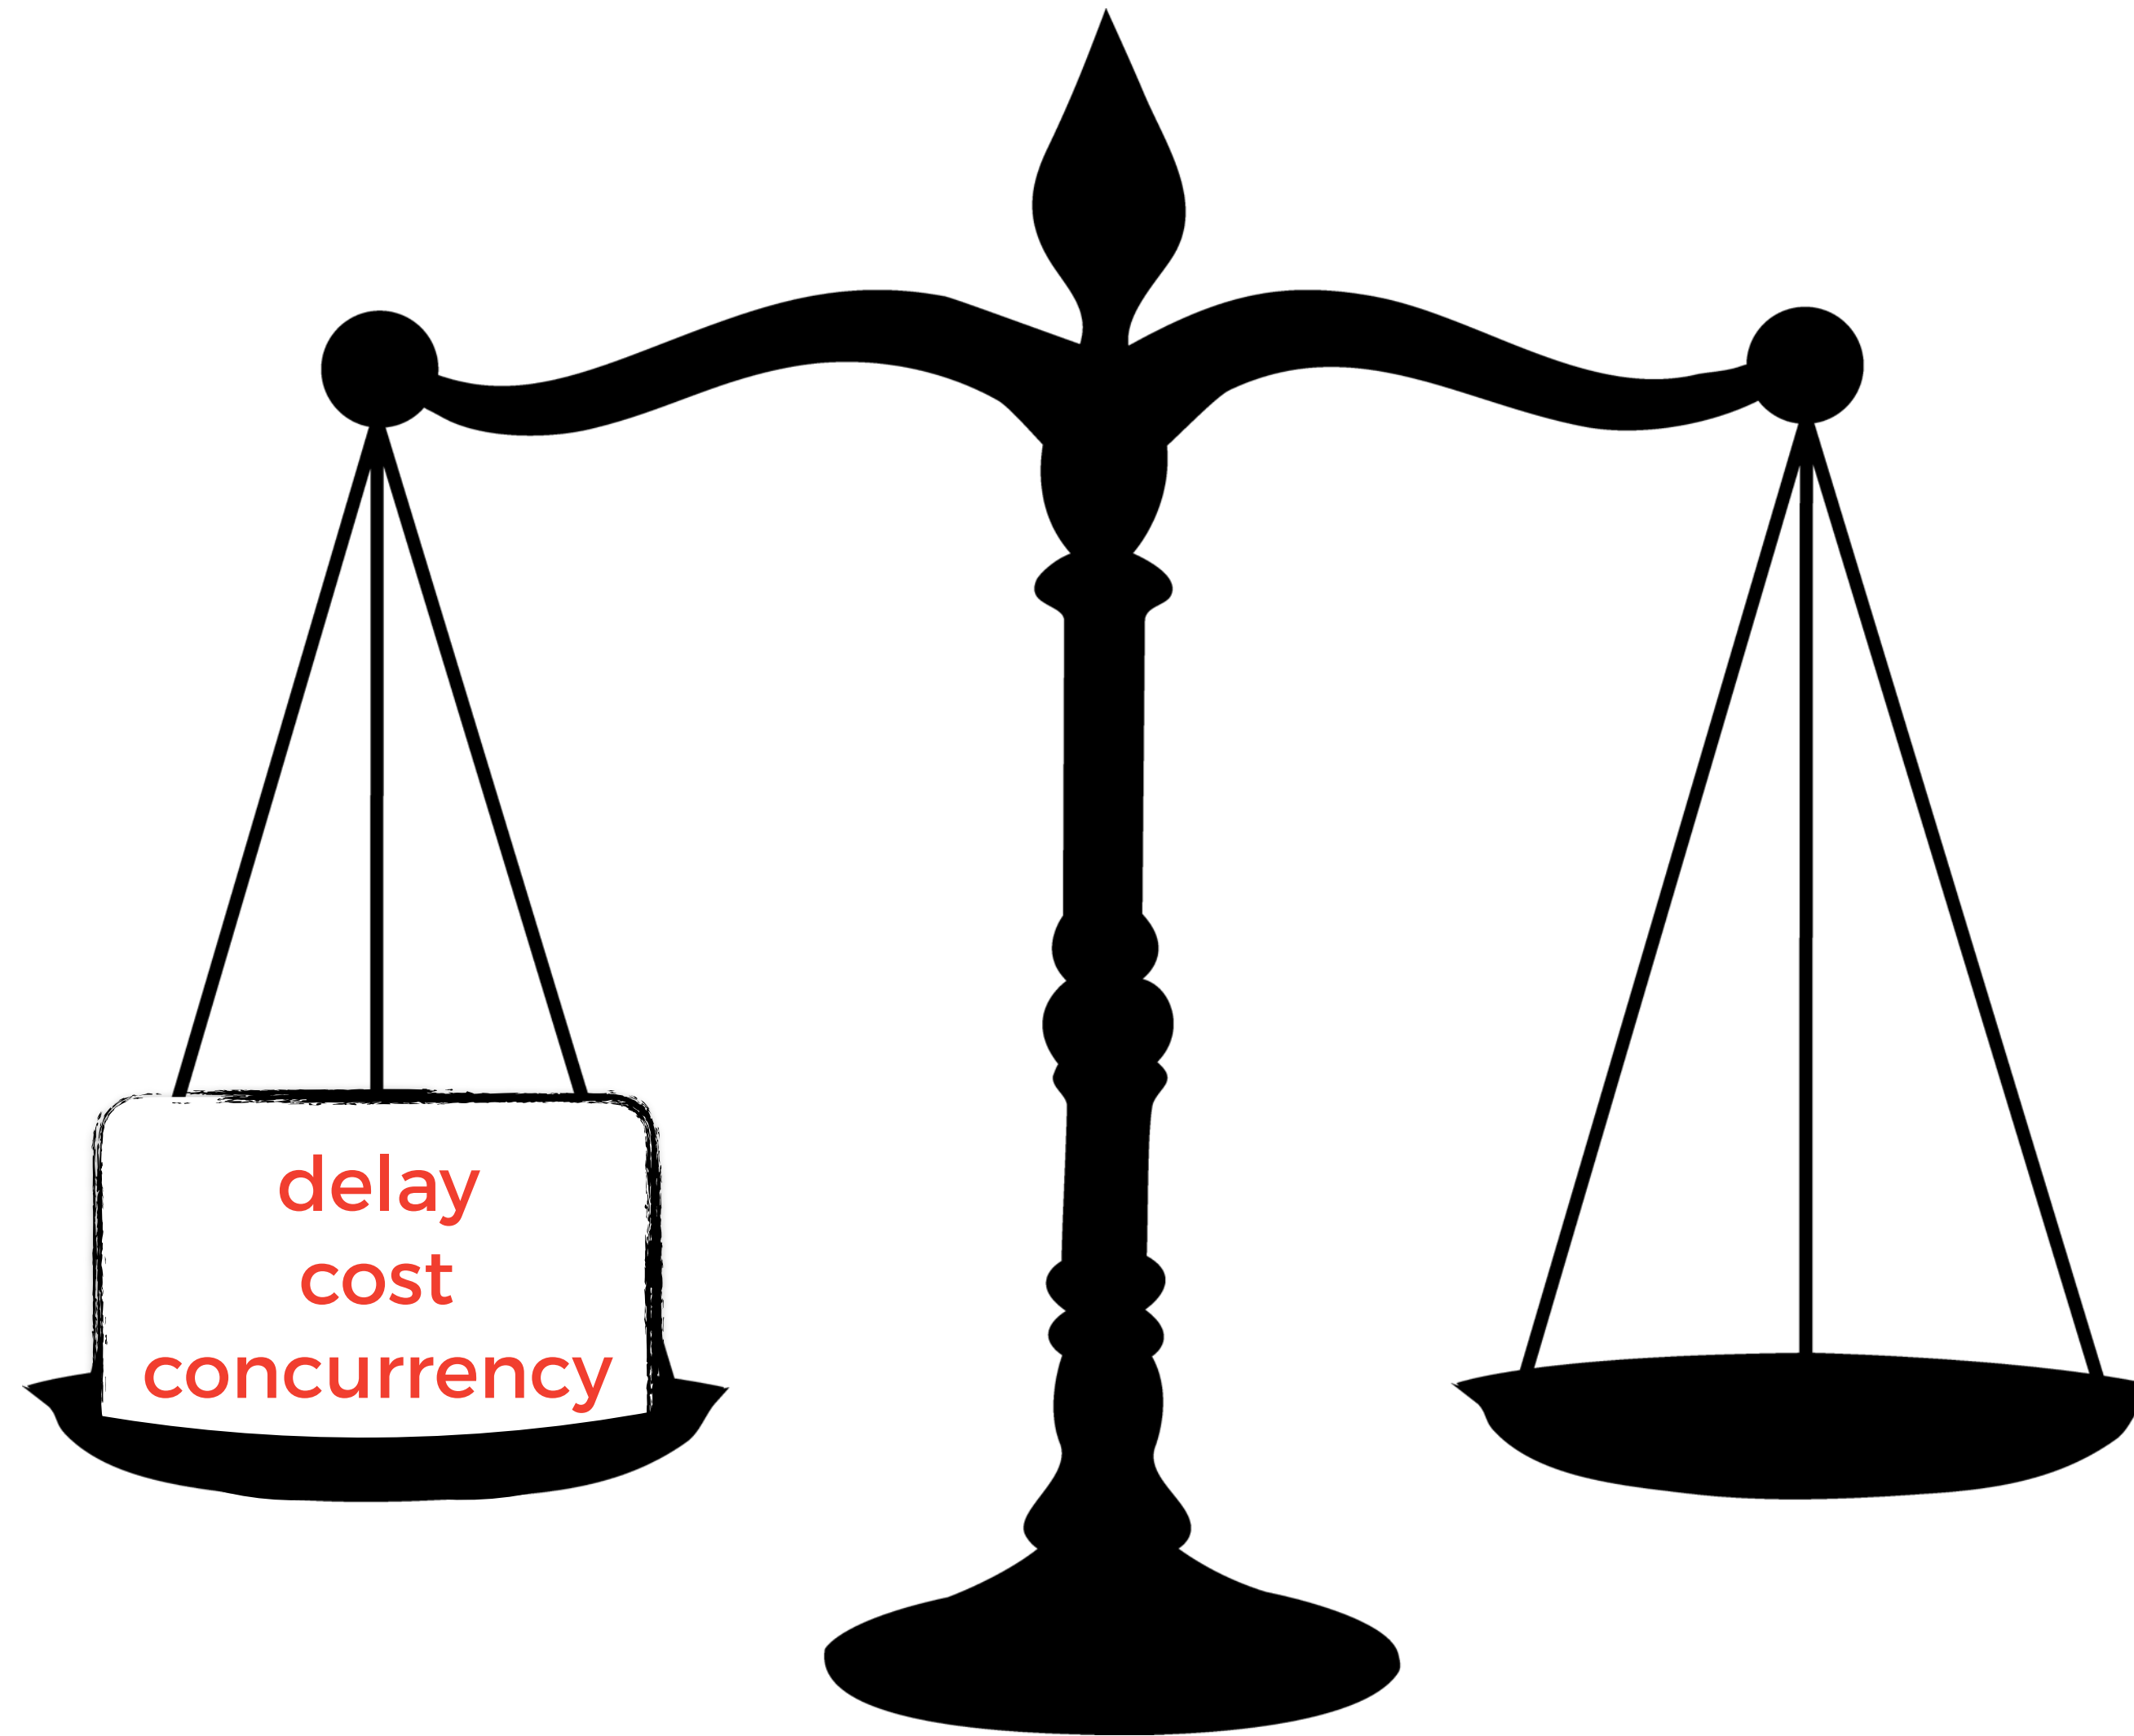




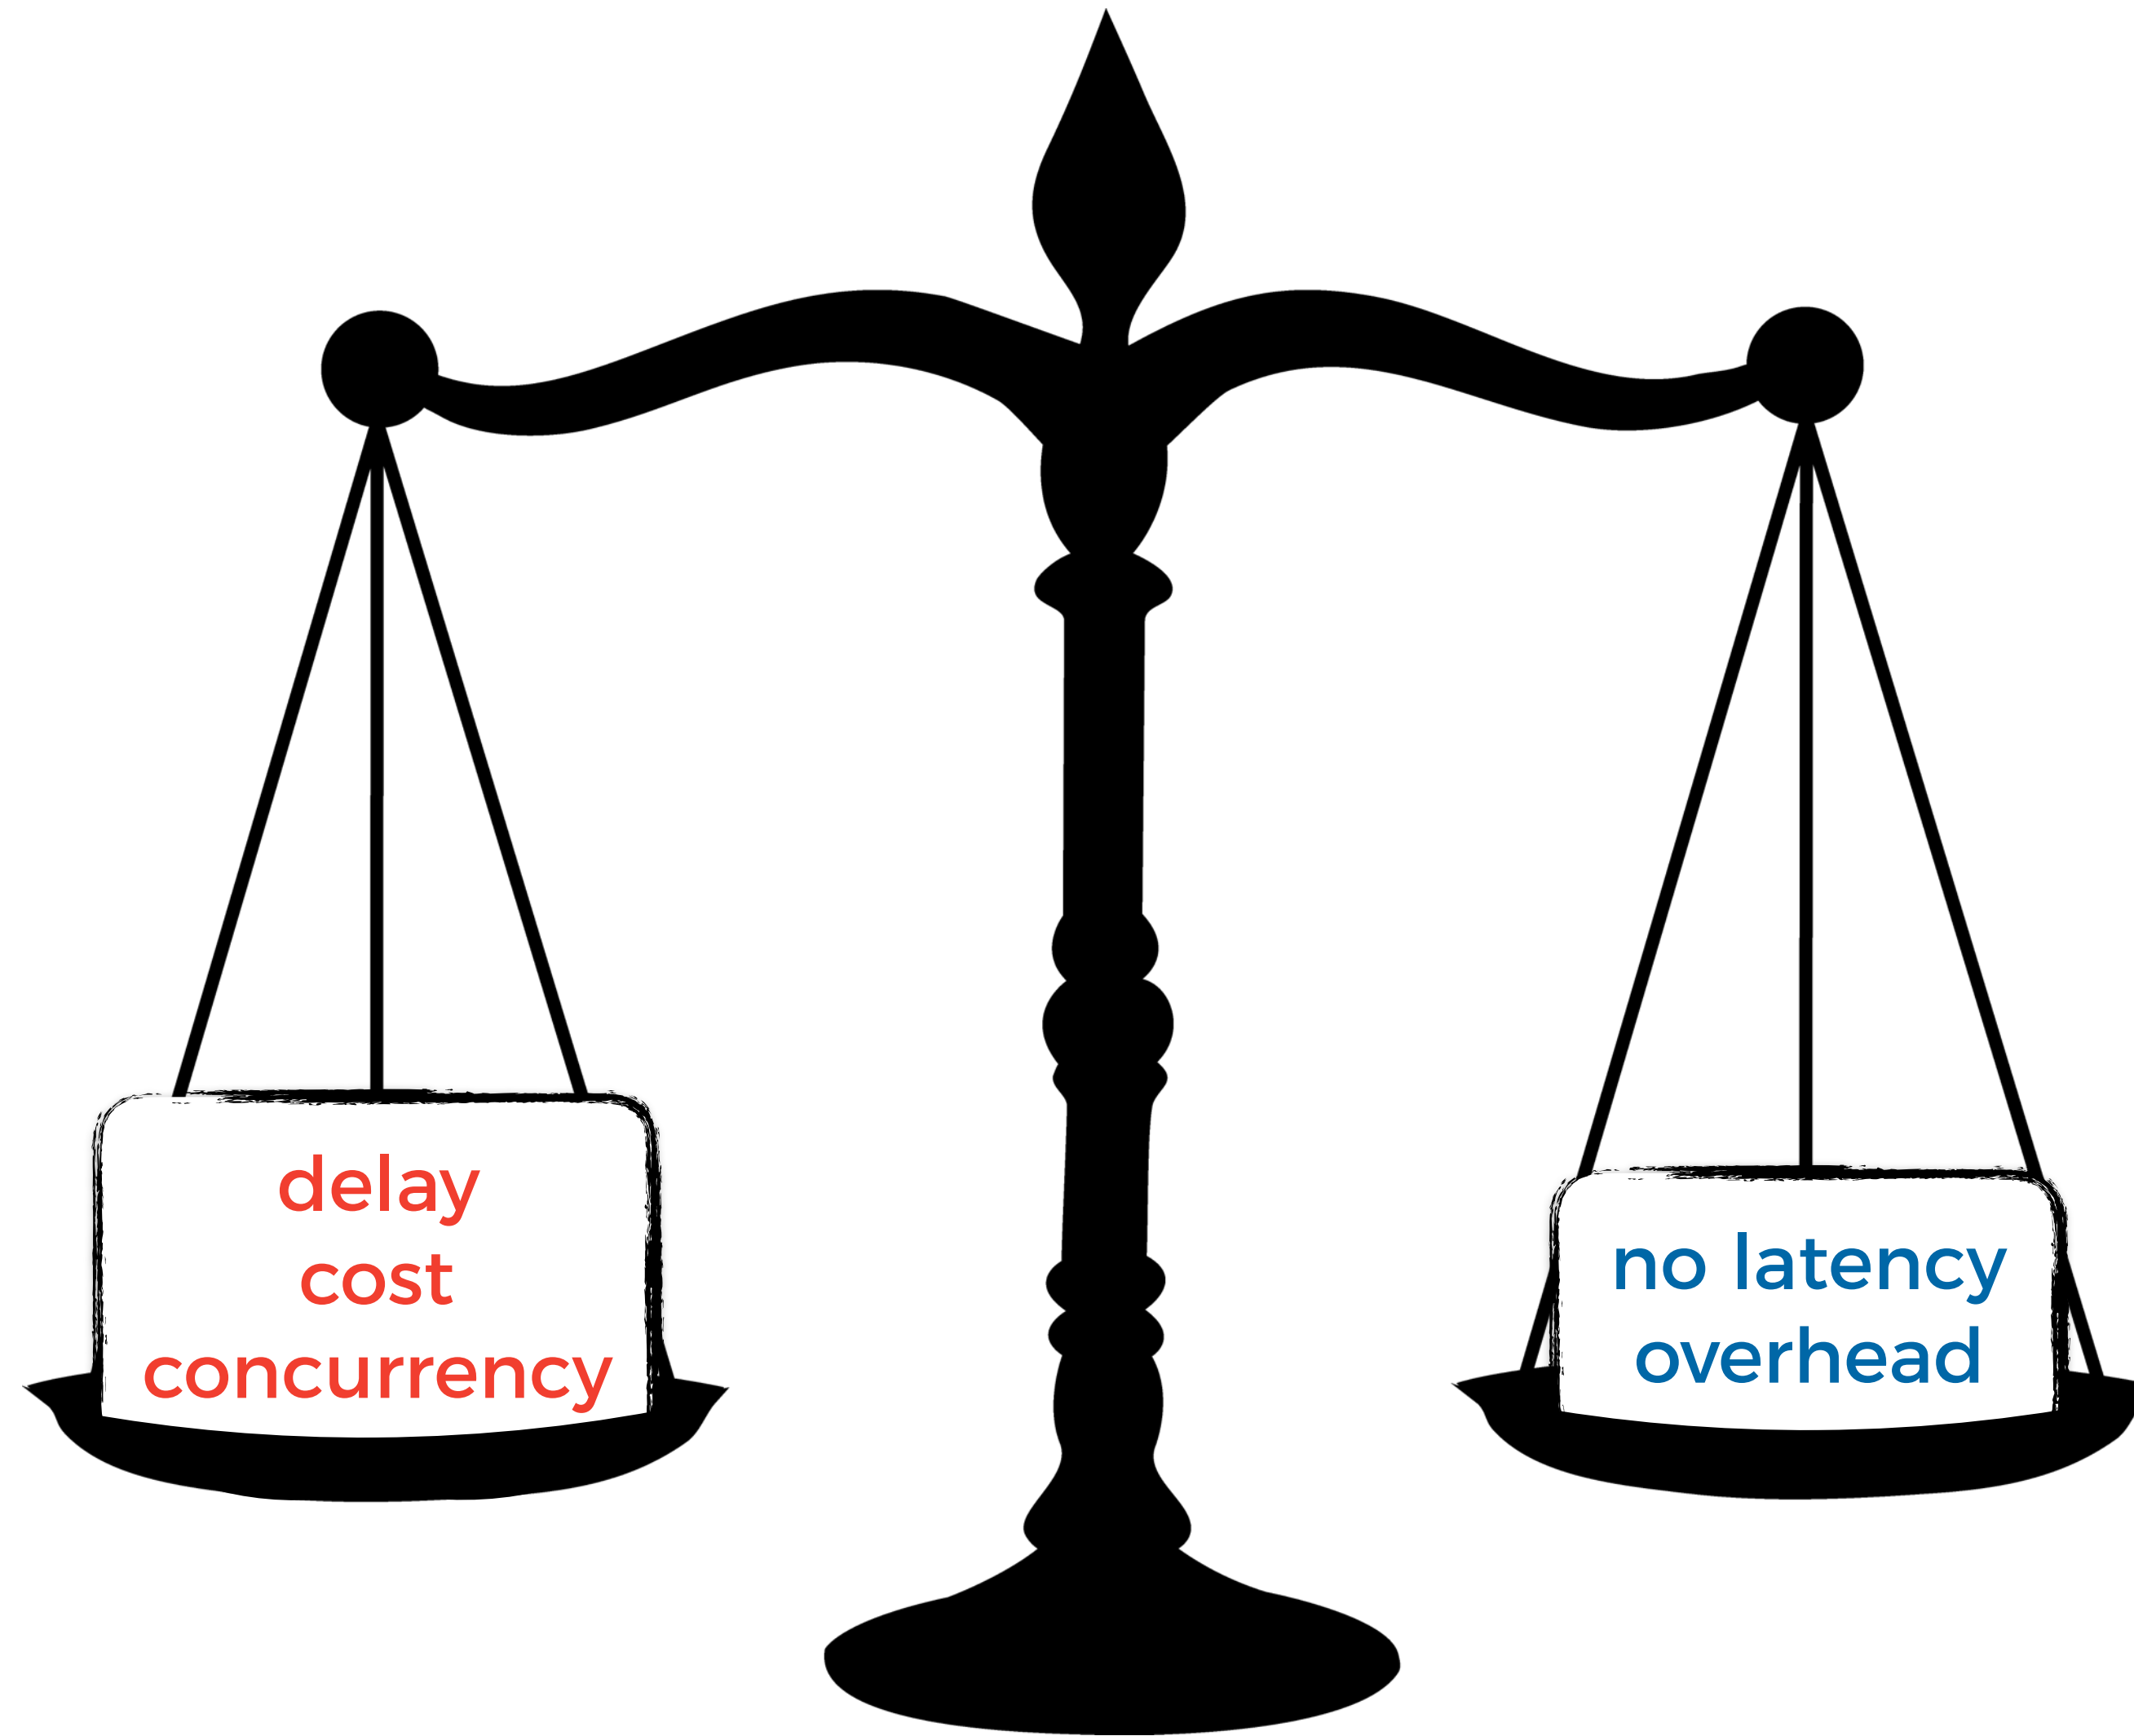
PURESEC



DASHBIRD



delay
cost
concurrency

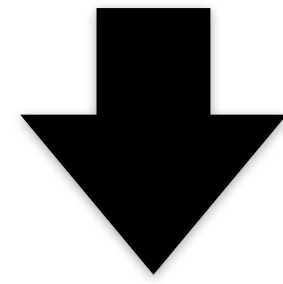


delay
cost
concurrency

no latency
overhead



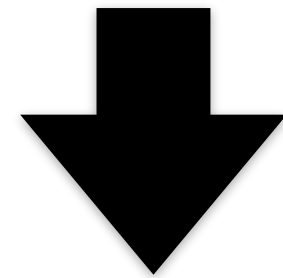
API Gateway



**send custom metrics
asynchronously**



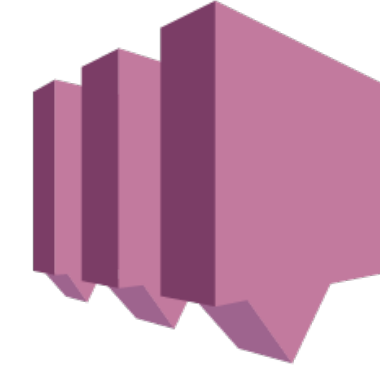
API Gateway



send custom metrics
asynchronously



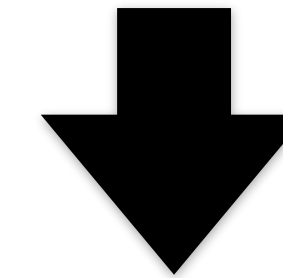
S3



SNS



Kinesis



send custom metrics as
part of function invocation

CloudWatch

Dashboards

* big-mouth

Alarms

ALARM

INSUFFICIENT

OK

Billing

Events

Rules

Event Buses

Logs

Metrics

Favorites

+ Add a dashboard

big-mouth

Add widget

Actions

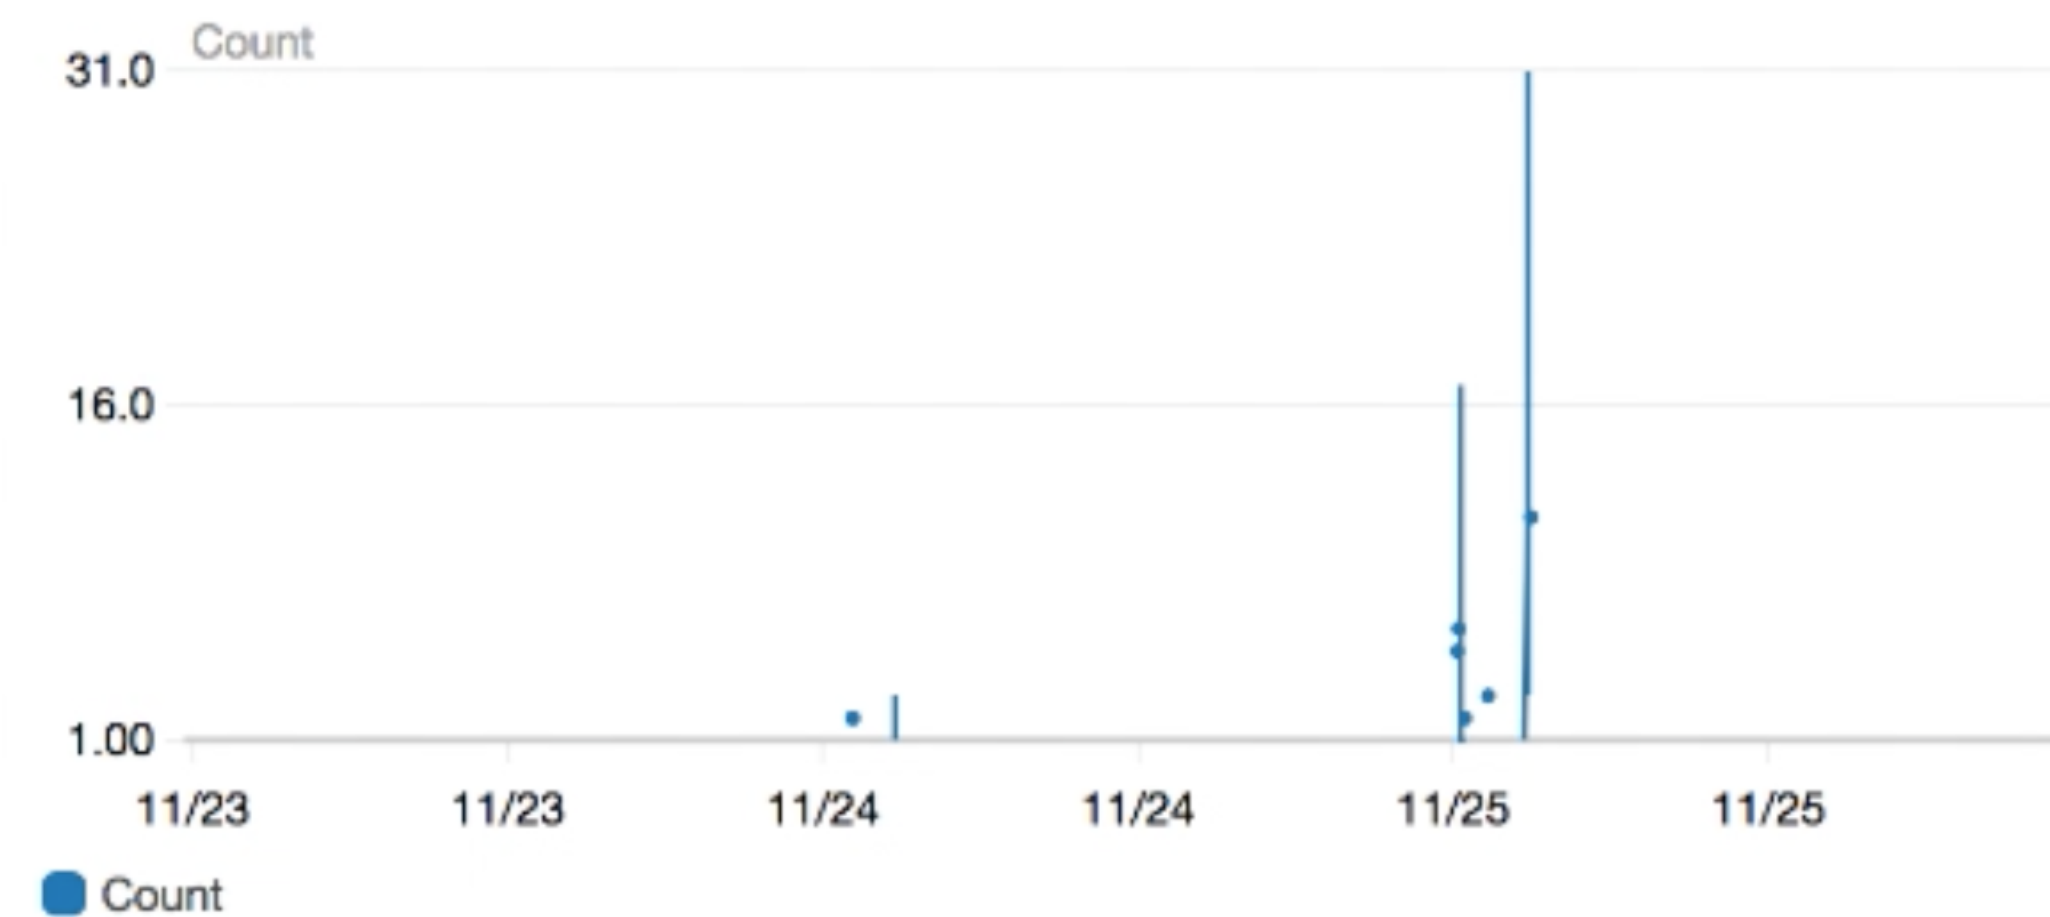
Save dashboard

1h 3h 12h 1d 3d 1w custom

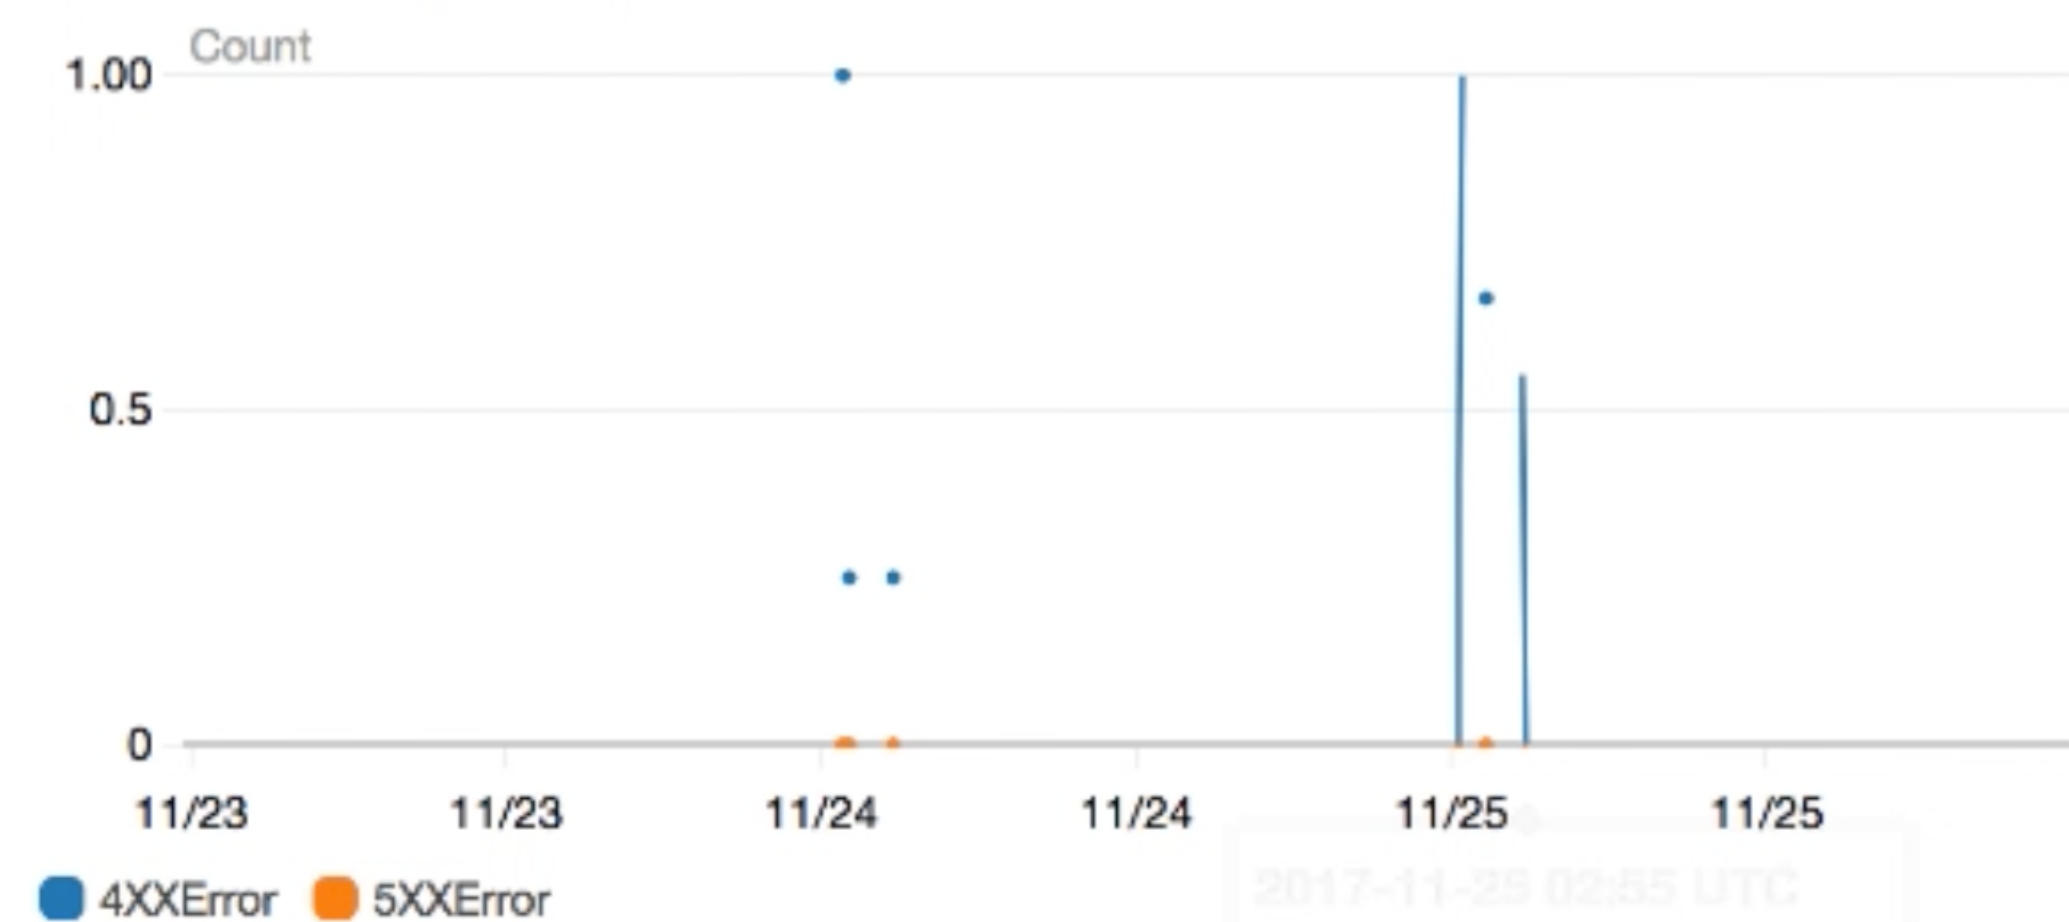
Latency



Count



4XXError, 5XXError



2017-11-25 02:55 UTC

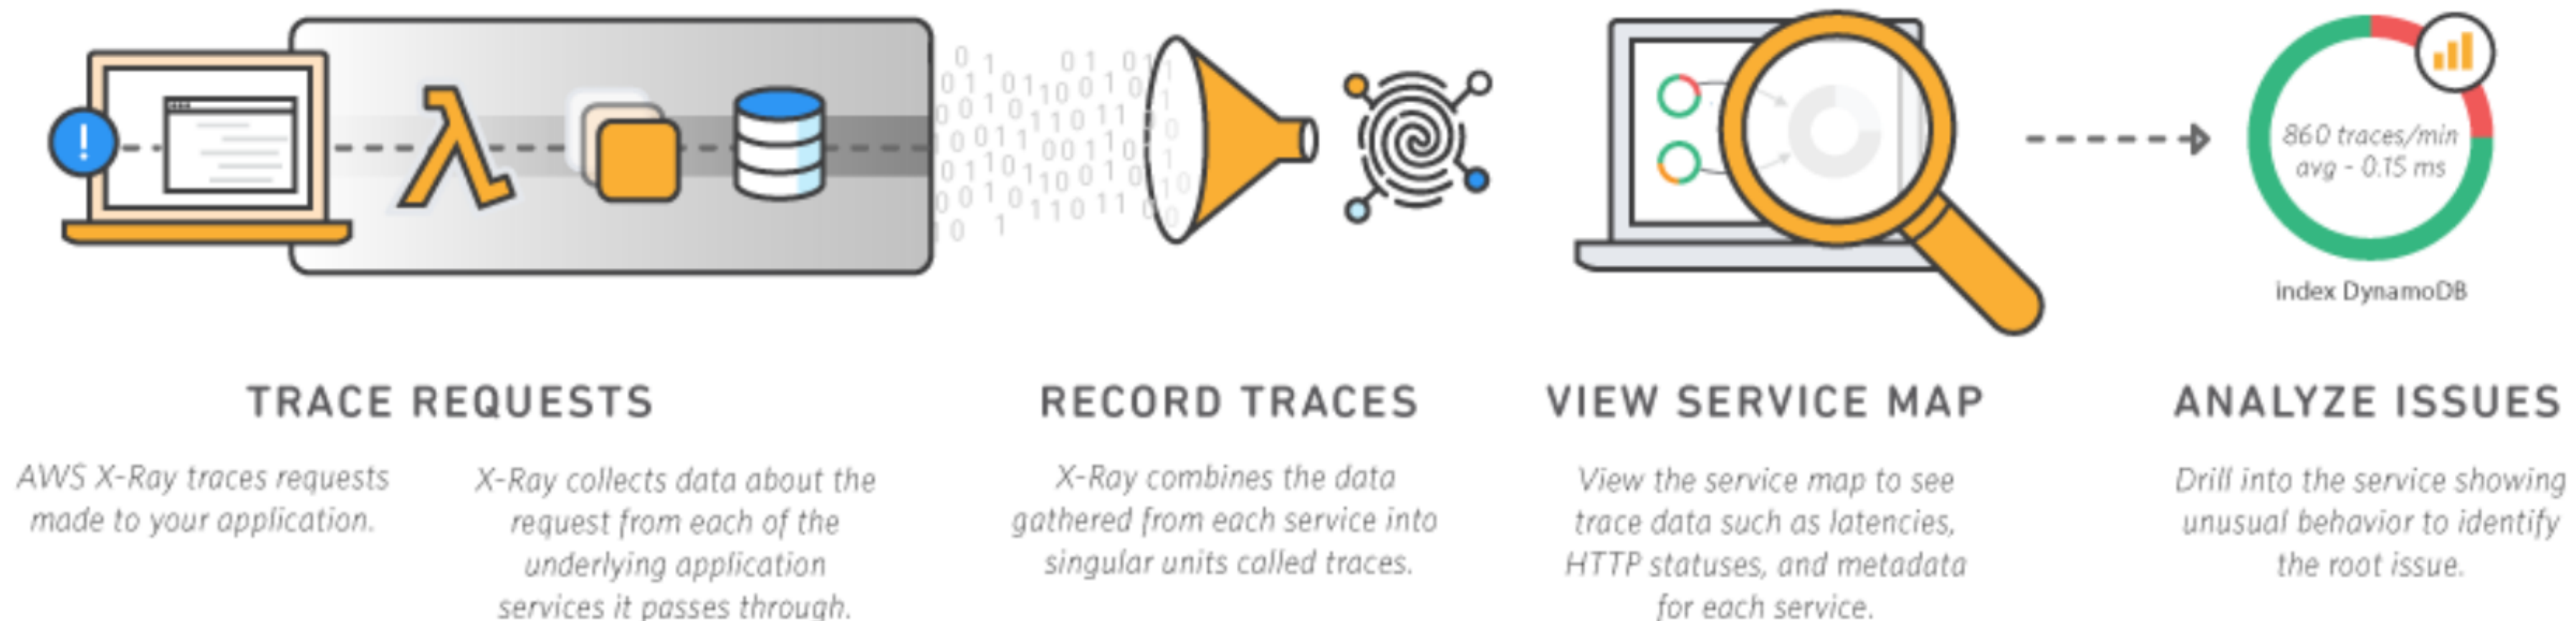
1. 4XXError

2. 5XXError



AWS X-Ray helps developers analyze and debug production, distributed applications, such as those built using a microservices architecture. With X-Ray, you can understand how your application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. X-Ray provides an end-to-end view of requests as they travel through your application, and shows a map of your application's underlying components. You can use X-Ray to analyze both applications in development and in production, from simple three-tier applications to complex microservices applications consisting of thousands of services.

How It Works



Name	Res.	Duration	Status	0.0ms	500ms	1.0s	1.5s	2.0s	2.5s	3.0s	3.5s
------	------	----------	--------	-------	-------	------	------	------	------	------	------

▼ **lambda-x-ray-demo-dev-service-a** AWS::Lambda

lambda-x-ray-demo-dev-service-a	200	3.2 sec	✓								
---------------------------------	-----	---------	---	--	--	--	--	--	--	--	--

▼ **lambda-x-ray-demo-dev-service-a** AWS::Lambda::Function

lambda-x-ray-demo-dev-service-a	-	2.1 sec	✓								
Initialization	-	529 ms	✓								
## publishing to SNS	-	158 ms	✓								
SNS	200	151 ms	✓								
Publish											
## accessing S3	-	110 ms	✓								
S3	404	65.0 ms	!								
GetObject											
S3	200	40.0 ms	✓								
PutObject											
## accessing DynamoDB	-	72.0 ms	✓								
DynamoDB	200	36.0 ms	✓								
GetItem: lambda-x-ray-demo-dev											
PutItem: lambda-x-ray-demo-dev											
## invoking service-c	-	1.2 sec	✓								
Lambda	200	1.2 sec	✓								
Invoke: lambda-x-ray-demo-dev-service-c											
## calling service b	-	550 ms	✓								
1g7cabdok5.execute-api.us-east-1.amazonav	200	548 ms	✓								
Remote: GET ... /dev/demo/service-b											

▼ **lambda-x-ray-demo-dev-service-c** AWS::Lambda

lambda-x-ray-demo-dev-service-c	200	1.1 sec	✓								
---------------------------------	-----	---------	---	--	--	--	--	--	--	--	--

▼ **lambda-x-ray-demo-dev-service-c** AWS::Lambda::Function

lambda-x-ray-demo-dev-service-c	-	230 ms	✓								
Initialization	-	443 ms	✓								
## publishing to SNS	-	193 ms	✓								
SNS	200	185 ms	✓								
Publish											

Name	Res.	Duration	Status	0.0ms	500ms	1.0s	1.5s	2.0s	2.5s	3.0s	3.5s
------	------	----------	--------	-------	-------	------	------	------	------	------	------

▼ **lambda-x-ray-demo-dev-service-a** AWS::Lambda

lambda-x-ray-demo-dev-service-a	200	3.2 sec	✓								
---------------------------------	-----	---------	---	--	--	--	--	--	--	--	--

▼ **lambda-x-ray-demo-dev-service-a** AWS::Lambda::Function

lambda-x-ray-demo-dev-service-a	-	2.1 sec	✓								
Initialization	-	529 ms	✓								
## publishing to SNS	-	158 ms	✓								
SNS	200	151 ms	✓								
Publish											
## accessing S3	-	110 ms	✓								
S3	404	65.0 ms	!								
GetObject											
S3	200	40.0 ms	✓								
PutObject											
## accessing DynamoDB	-	72.0 ms	✓								
DynamoDB	200	36.0 ms	✓								
GetItem: lambda-x-ray-demo-dev											
DynamoDB	200	32.0 ms	✓								
PutItem: lambda-x-ray-demo-dev											
## invoking service-c	-	1.2 sec	✓								
Lambda	200	1.2 sec	✓								
Invoke: lambda-x-ray-demo-dev-service-c											
## calling service b	-	550 ms	✓								
1g7cabdok5.execute-api.us-east-1.amazonaws.com	200	548 ms	✓								
Remote: GET ... /dev/demo/service-b											



▼ **lambda-x-ray-demo-dev-service-c** AWS::Lambda

lambda-x-ray-demo-dev-service-c	200	1.1 sec	✓								
---------------------------------	-----	---------	---	--	--	--	--	--	--	--	--

▼ **lambda-x-ray-demo-dev-service-c** AWS::Lambda::Function

lambda-x-ray-demo-dev-service-c	-	230 ms	✓								
Initialization	-	443 ms	✓								
## publishing to SNS	-	193 ms	✓								
SNS	200	185 ms	✓								
Publish											

Name	Res.	Duration	Status	0.0ms	500ms	1.0s	1.5s	2.0s	2.5s	3.0s	3.5s
------	------	----------	--------	-------	-------	------	------	------	------	------	------

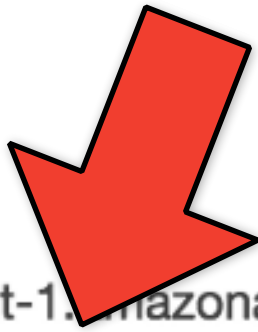
▼ lambda-x-ray-demo-dev-service-a AWS::Lambda

lambda-x-ray-demo-dev-service-a	200	3.2 sec	✓	[Timeline bar from 0.0ms to 3.2s]							
---------------------------------	-----	---------	---	-----------------------------------	--	--	--	--	--	--	--

▼ lambda-x-ray-demo-dev-service-a AWS::Lambda::Function

lambda-x-ray-demo-dev-service-a	-	2.1 sec	✓	[Timeline bar from 0.0ms to 2.1s]							
Initialization	-	529 ms	✓	[Timeline bar from 0.5s to 1.0s]							
## publishing to SNS	-	158 ms	✓	[Timeline bar from 1.0s to 1.15s]							
SNS	200	151 ms	✓	[Timeline bar from 1.0s to 1.15s]							
## accessing S3	-	110 ms	✓	[Timeline bar from 1.15s to 1.26s]							
S3	404	65.0 ms	!	[Timeline bar from 1.15s to 1.26s]							
S3	-	-	-	[Timeline bar from 1.15s to 1.26s]							
## accessing DynamoDB	-	-	-	[Timeline bar from 1.15s to 1.26s]							
DynamoDB	-	-	-	[Timeline bar from 1.15s to 1.26s]							
DynamoDB	200	32.0 ms	✓	[Timeline bar from 1.15s to 1.26s]							
## invoking service-c	-	1.2 sec	✓	[Timeline bar from 1.26s to 2.46s]							
Lambda	200	1.2 sec	✓	[Timeline bar from 1.26s to 2.46s]							
## calling service b	-	550 ms	✓	[Timeline bar from 2.46s to 3.01s]							
1g7cabdok5.execute-api.us-east-1.amazonaws.com	200	548 ms	✓	[Timeline bar from 2.46s to 3.01s]							

do not span over API Gateway



▼ lambda-x-ray-demo-dev-service-c AWS::Lambda

lambda-x-ray-demo-dev-service-c	200	1.1 sec	✓	[Timeline bar from 1.5s to 2.6s]							
---------------------------------	-----	---------	---	----------------------------------	--	--	--	--	--	--	--

▼ lambda-x-ray-demo-dev-service-c AWS::Lambda::Function

lambda-x-ray-demo-dev-service-c	-	230 ms	✓	[Timeline bar from 1.5s to 1.73s]							
Initialization	-	443 ms	✓	[Timeline bar from 1.0s to 1.44s]							
## publishing to SNS	-	193 ms	✓	[Timeline bar from 1.44s to 1.63s]							
SNS	200	185 ms	✓	[Timeline bar from 1.44s to 1.63s]							

Name	Res.	Duration	Status	0.0ms	500ms	1.0s	1.5s	2.0s	2.5s	3.0s	3.5s
------	------	----------	--------	-------	-------	------	------	------	------	------	------

▼ lambda-x-ray-demo-dev-service-a AWS::Lambda

lambda-x-ray-demo-dev-service-a	200	3.2 sec	✓	[Timeline bar]							
---------------------------------	-----	---------	---	----------------	--	--	--	--	--	--	--

▼ lambda-x-ray-demo-dev-service-a AWS::Lambda::Function

lambda-x-ray-demo-dev-service-a	-	2.1 sec	✓	[Timeline bar]							
---------------------------------	---	---------	---	----------------	--	--	--	--	--	--	--

- Initialization
- ## publishing to SNS
 - SNS
- ## accessing S3
 - S3
- S3
- ## accessing DynamoDB
 - DynamoDB
 - DynamoDB
- ## invoking service-c
 - Lambda
- ## calling service b
 - 1g7cabdok5.execute-api.us-east-1.amazonaws.com

narrow focus on a function

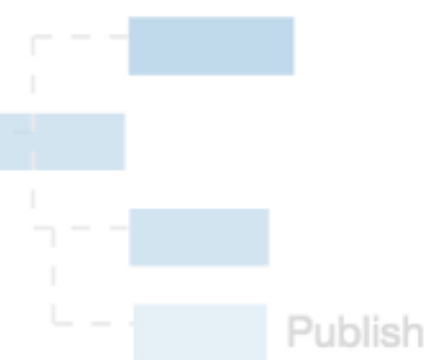
good for homing in on performance issues for a particular function, but offers little to help you build **intuition** about how your system operates as a whole.

▼ lambda-x-ray-demo-dev-service-c AWS::Lambda

lambda-x-ray-demo-dev-service-c	200	1.1 sec	✓	[Timeline bar]							
---------------------------------	-----	---------	---	----------------	--	--	--	--	--	--	--

▼ lambda-x-ray-demo-dev-service-c AWS::Lambda::Function

lambda-x-ray-demo-dev-service-c	-	230 ms	✓	[Timeline bar]							
Initialization	-	443 ms	✓	[Timeline bar]							
## publishing to SNS	-	193 ms	✓	[Timeline bar]							
SNS	200	185 ms	✓	[Timeline bar]							



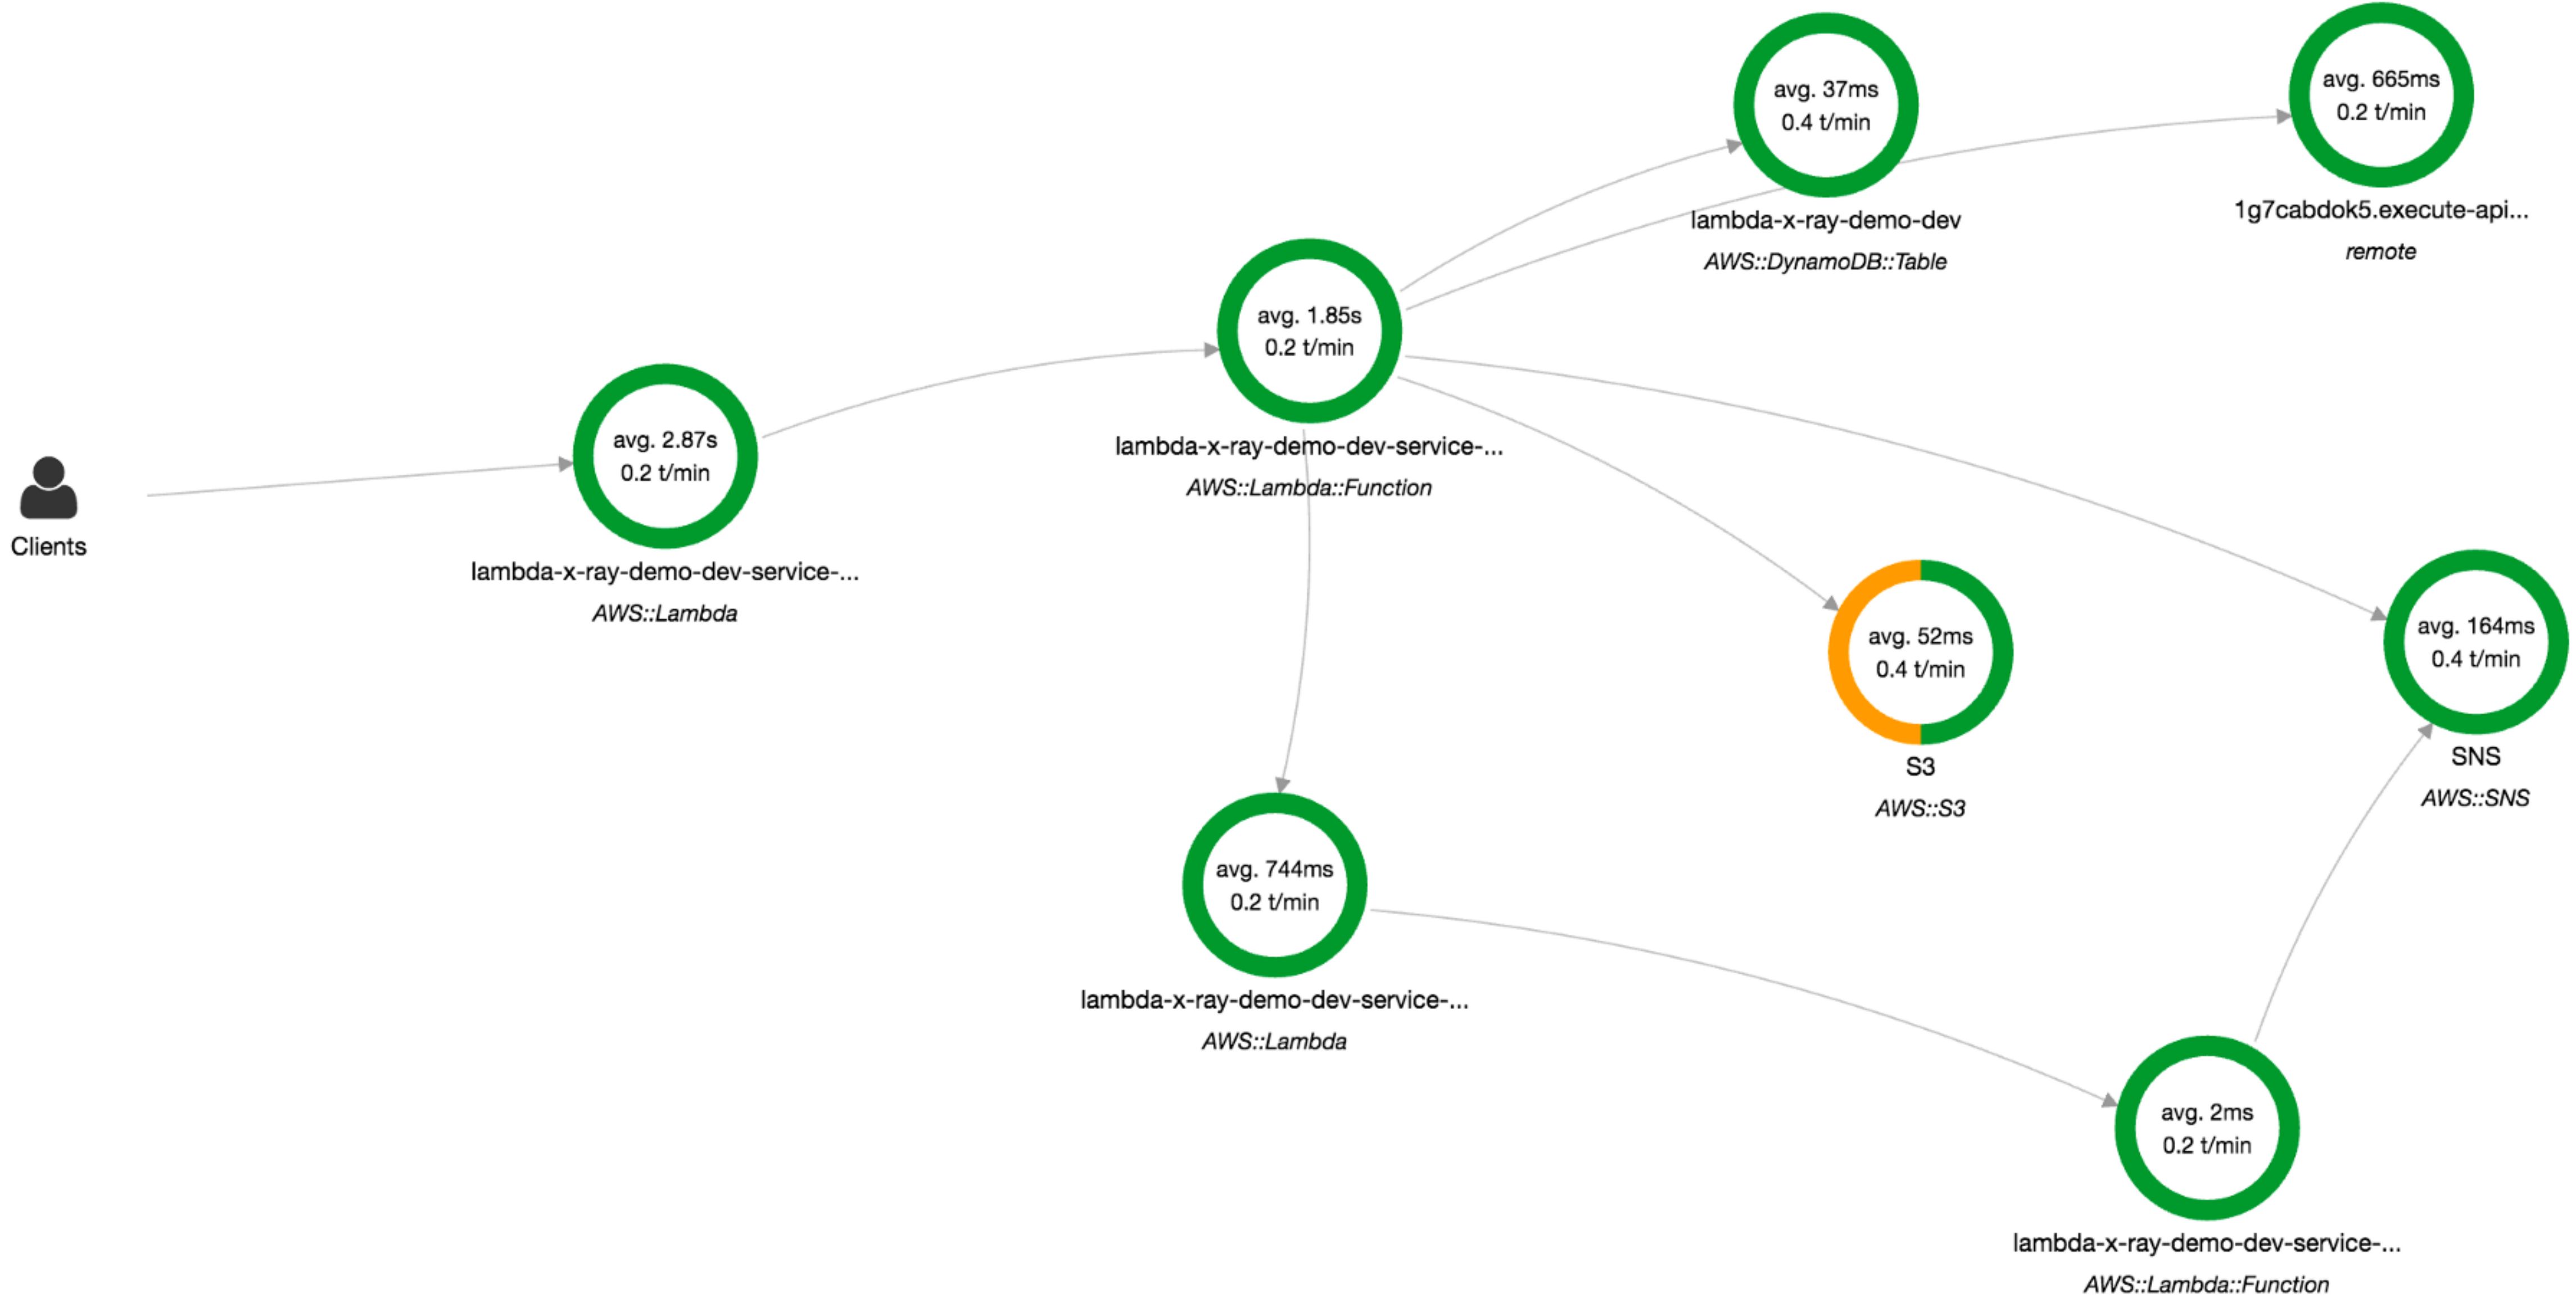


*However, I would argue that the health of the system no longer matters. We've entered an era where what matters is the health of each individual event, or each individual user's experience, or each shopping cart's experience (or other high cardinality dimensions). With distributed systems you don't care about the health of the system, you care about the **health of the event or the slice.** ””*

- Charity Majors <http://bit.ly/2E2QngU>

A photograph of Sherlock Holmes, played by Benedict Cumberbatch, pointing his right index finger towards the left. He is wearing a dark blue turtleneck sweater and a dark grey coat. A large, fluffy brown dog is in the foreground, looking towards the left. A white speech bubble with a black outline is positioned to the left of the dog, containing the text "follow the data". The background shows a white building with arched windows and a black metal fence.

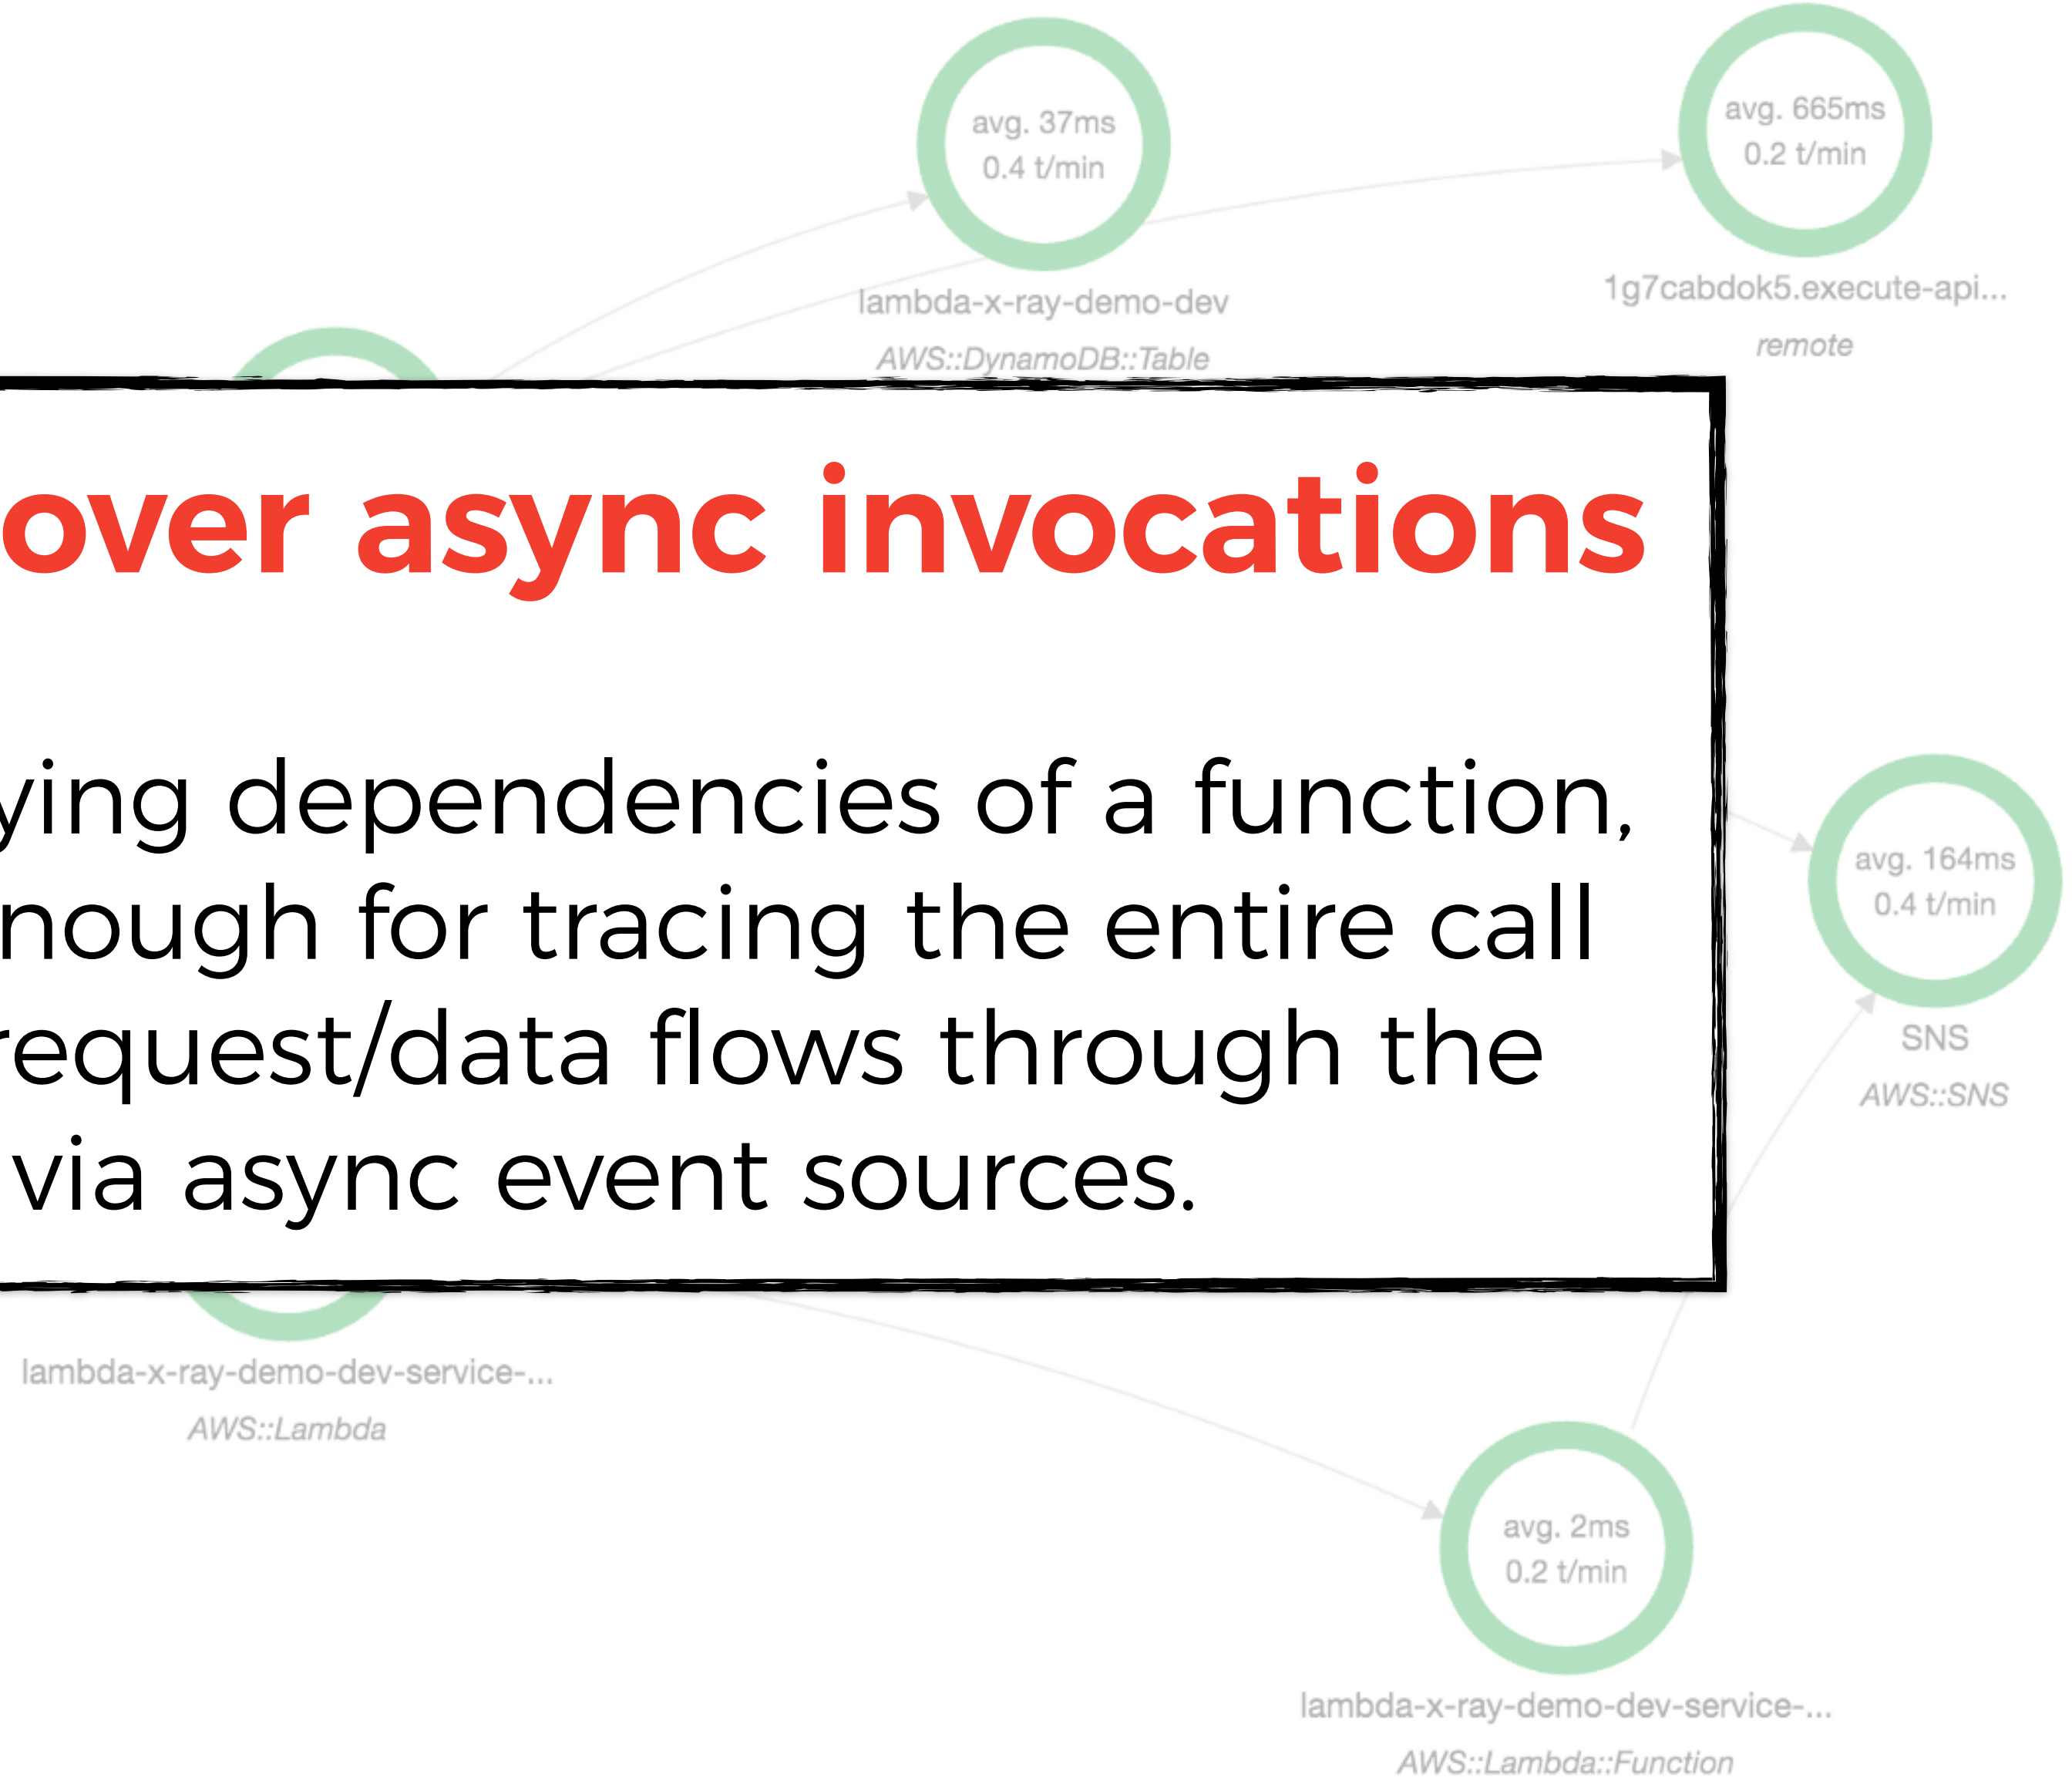
follow the data





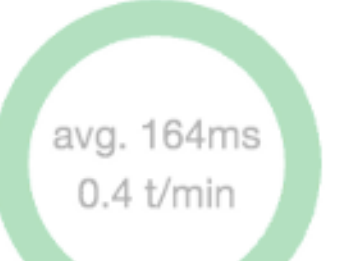
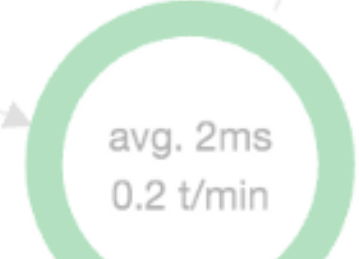
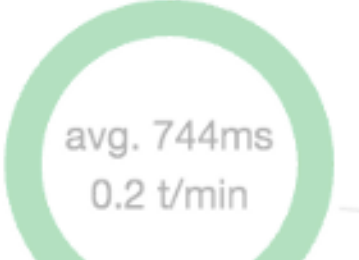
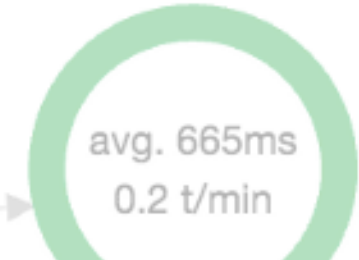
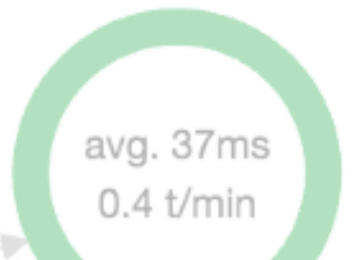
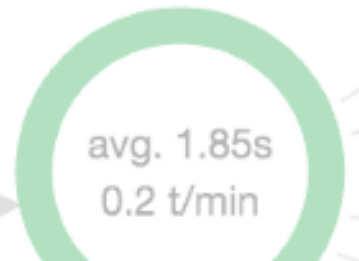
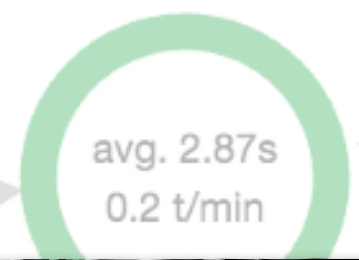
don't span over async invocations

good for identifying dependencies of a function, but not good enough for tracing the entire call chain as user request/data flows through the system via async event sources.





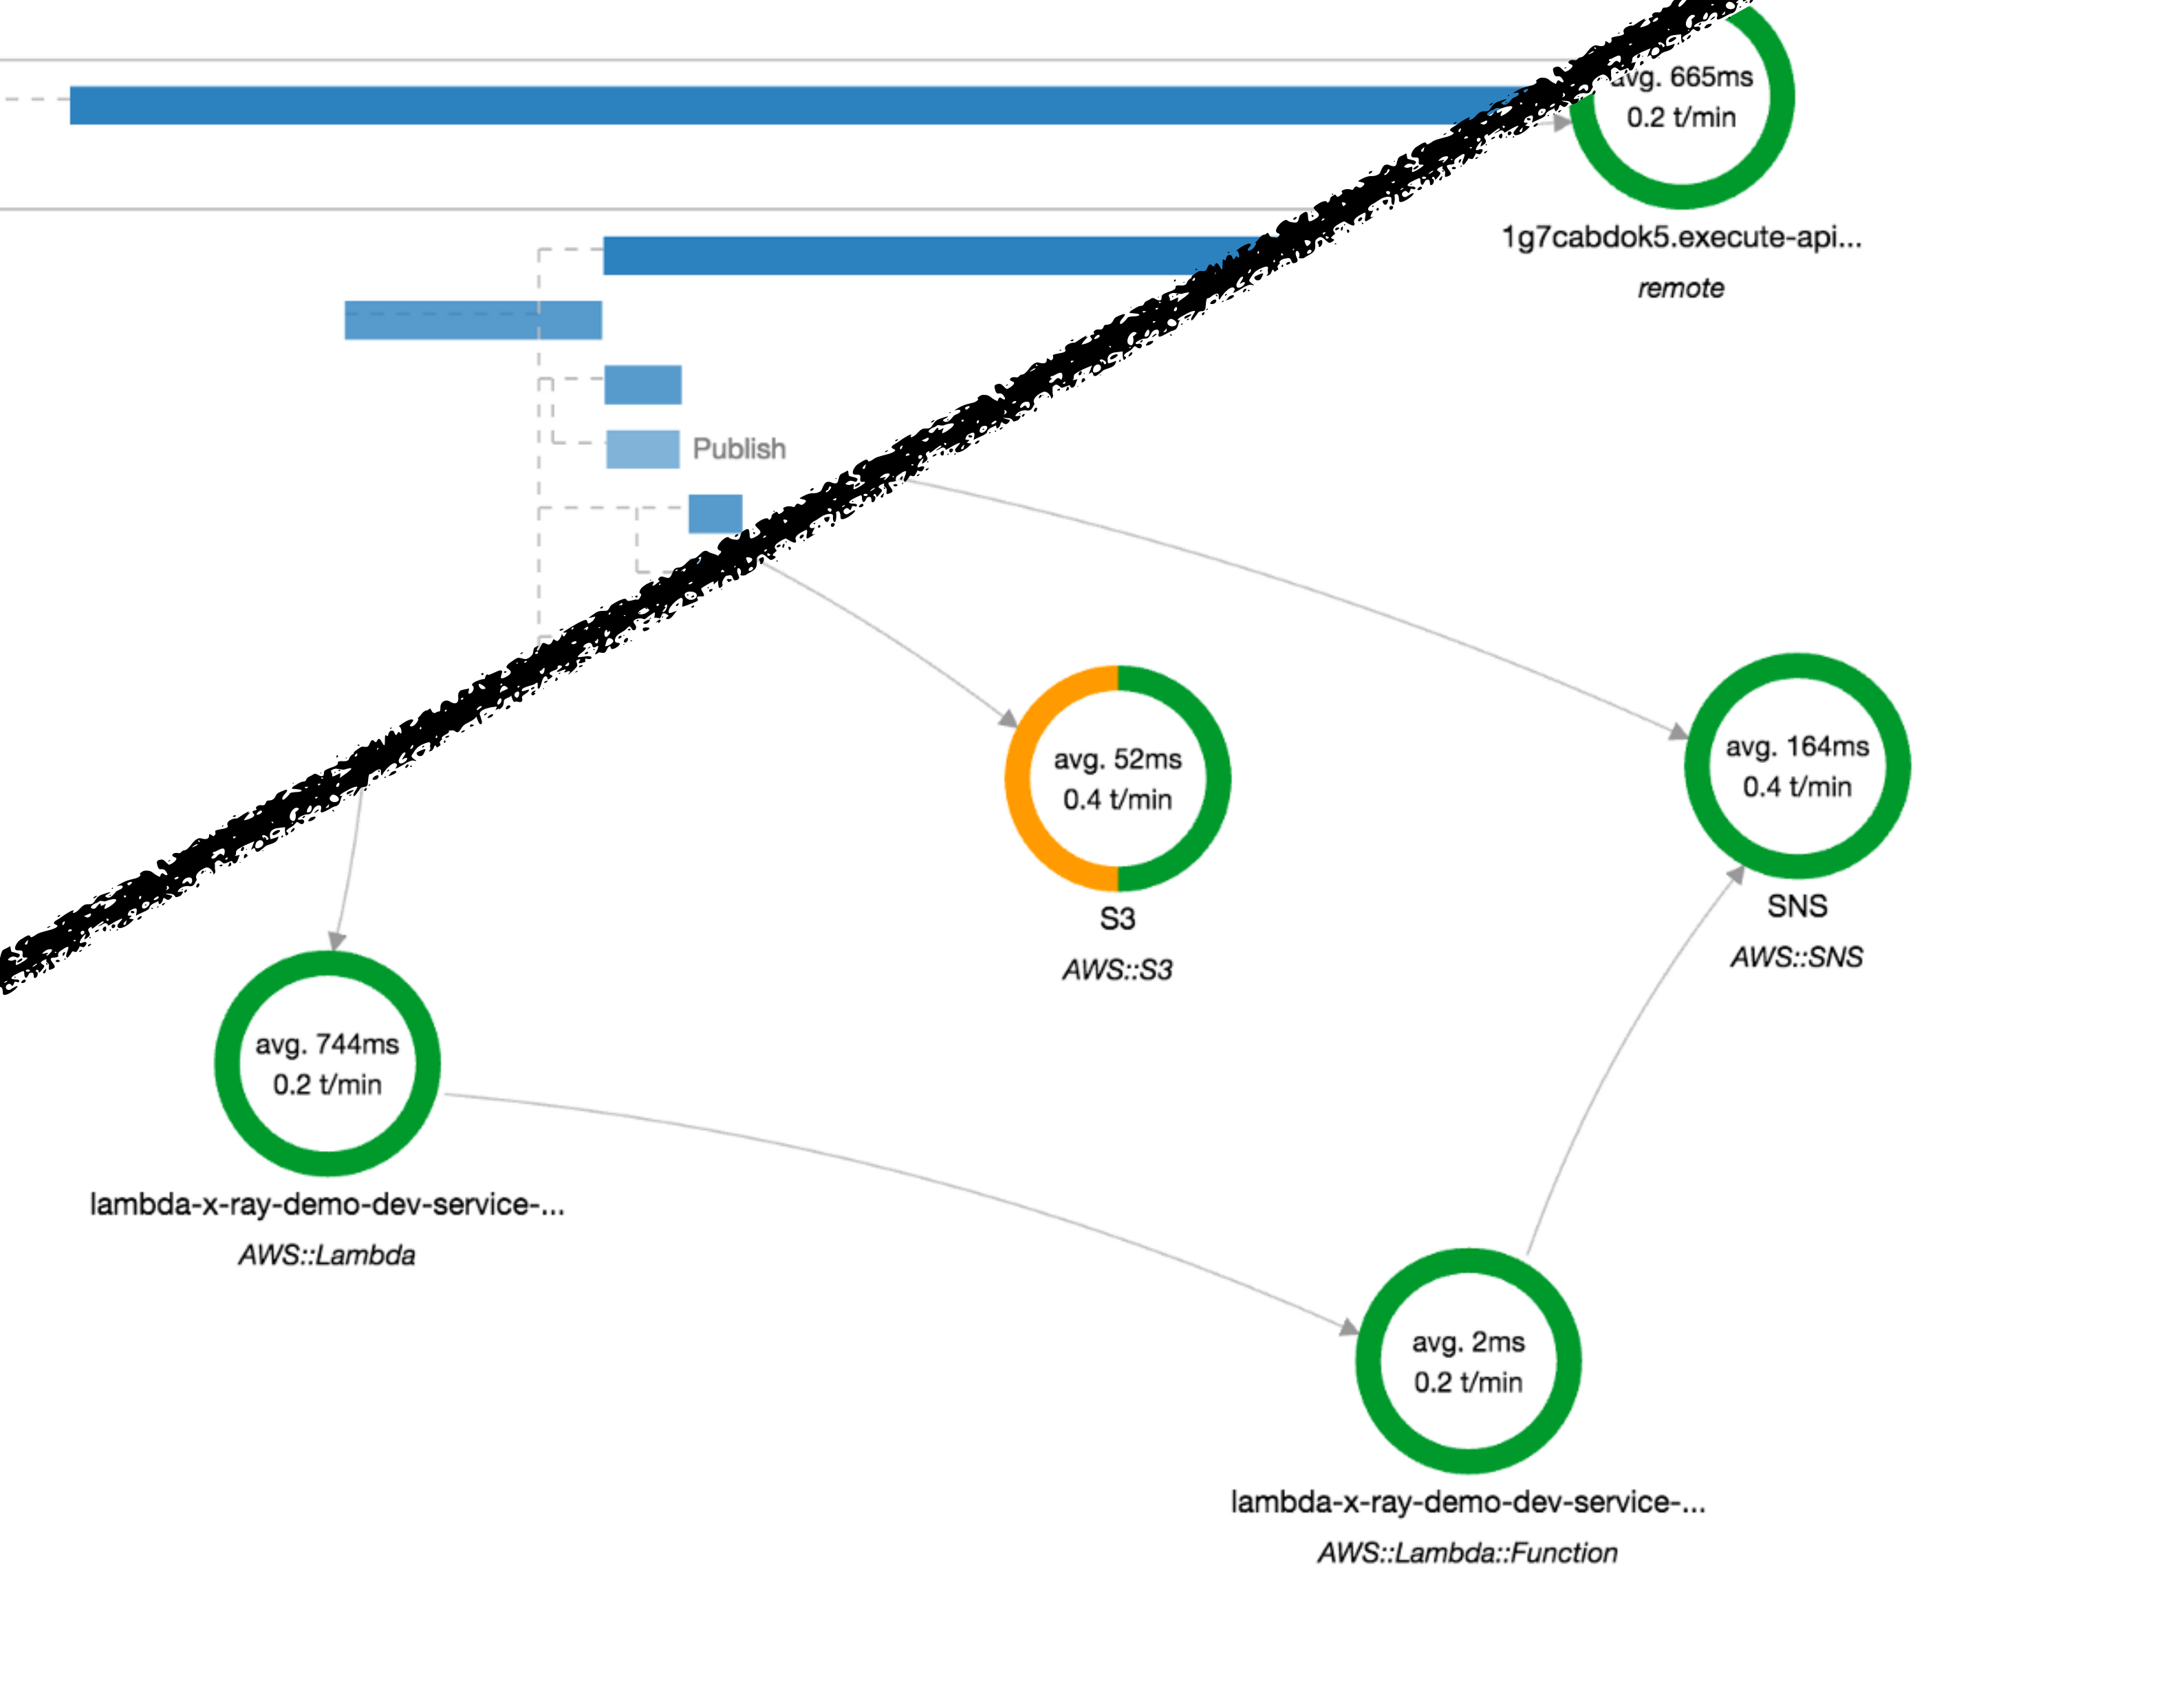
don't span over non-AWS services



Name	Res.	Duration	Status	0.0ms	500ms	1.0s	1.5s	2.0s	2.5s	3.0s	3.5s	4.0s
------	------	----------	--------	-------	-------	------	------	------	------	------	------	------

▼ lambda-x-ray-demo-dev-service-a AWS::Lambda			
lambda-x-ray-demo-dev-service-a	200	3.2 sec	✓
▼ lambda-x-ray-demo-dev-service-a AWS::Lambda::Function			
lambda-x-ray-demo-dev-service-a	-	2.1 sec	✓
Initialization	-	529 ms	✓
## publishing to SNS	-	158 ms	✓
SNS	200	151 ms	✓
## accessing S3	-	110 ms	✓
S3	404	65.0 ms	!
S3	200	40.0 ms	✓
## accessing DynamoDB	-	72.0 ms	✓
DynamoDB	200	36.0 ms	✓
DynamoDB	200	32.0 ms	✓
## invoking service-c	-	1.2 sec	✓
Lambda	200	1.2 sec	✓
## calling service b	-	550 ms	✓
1g7cabdok5.execute-api.us-east-1.amazonav	200	548 ms	✓

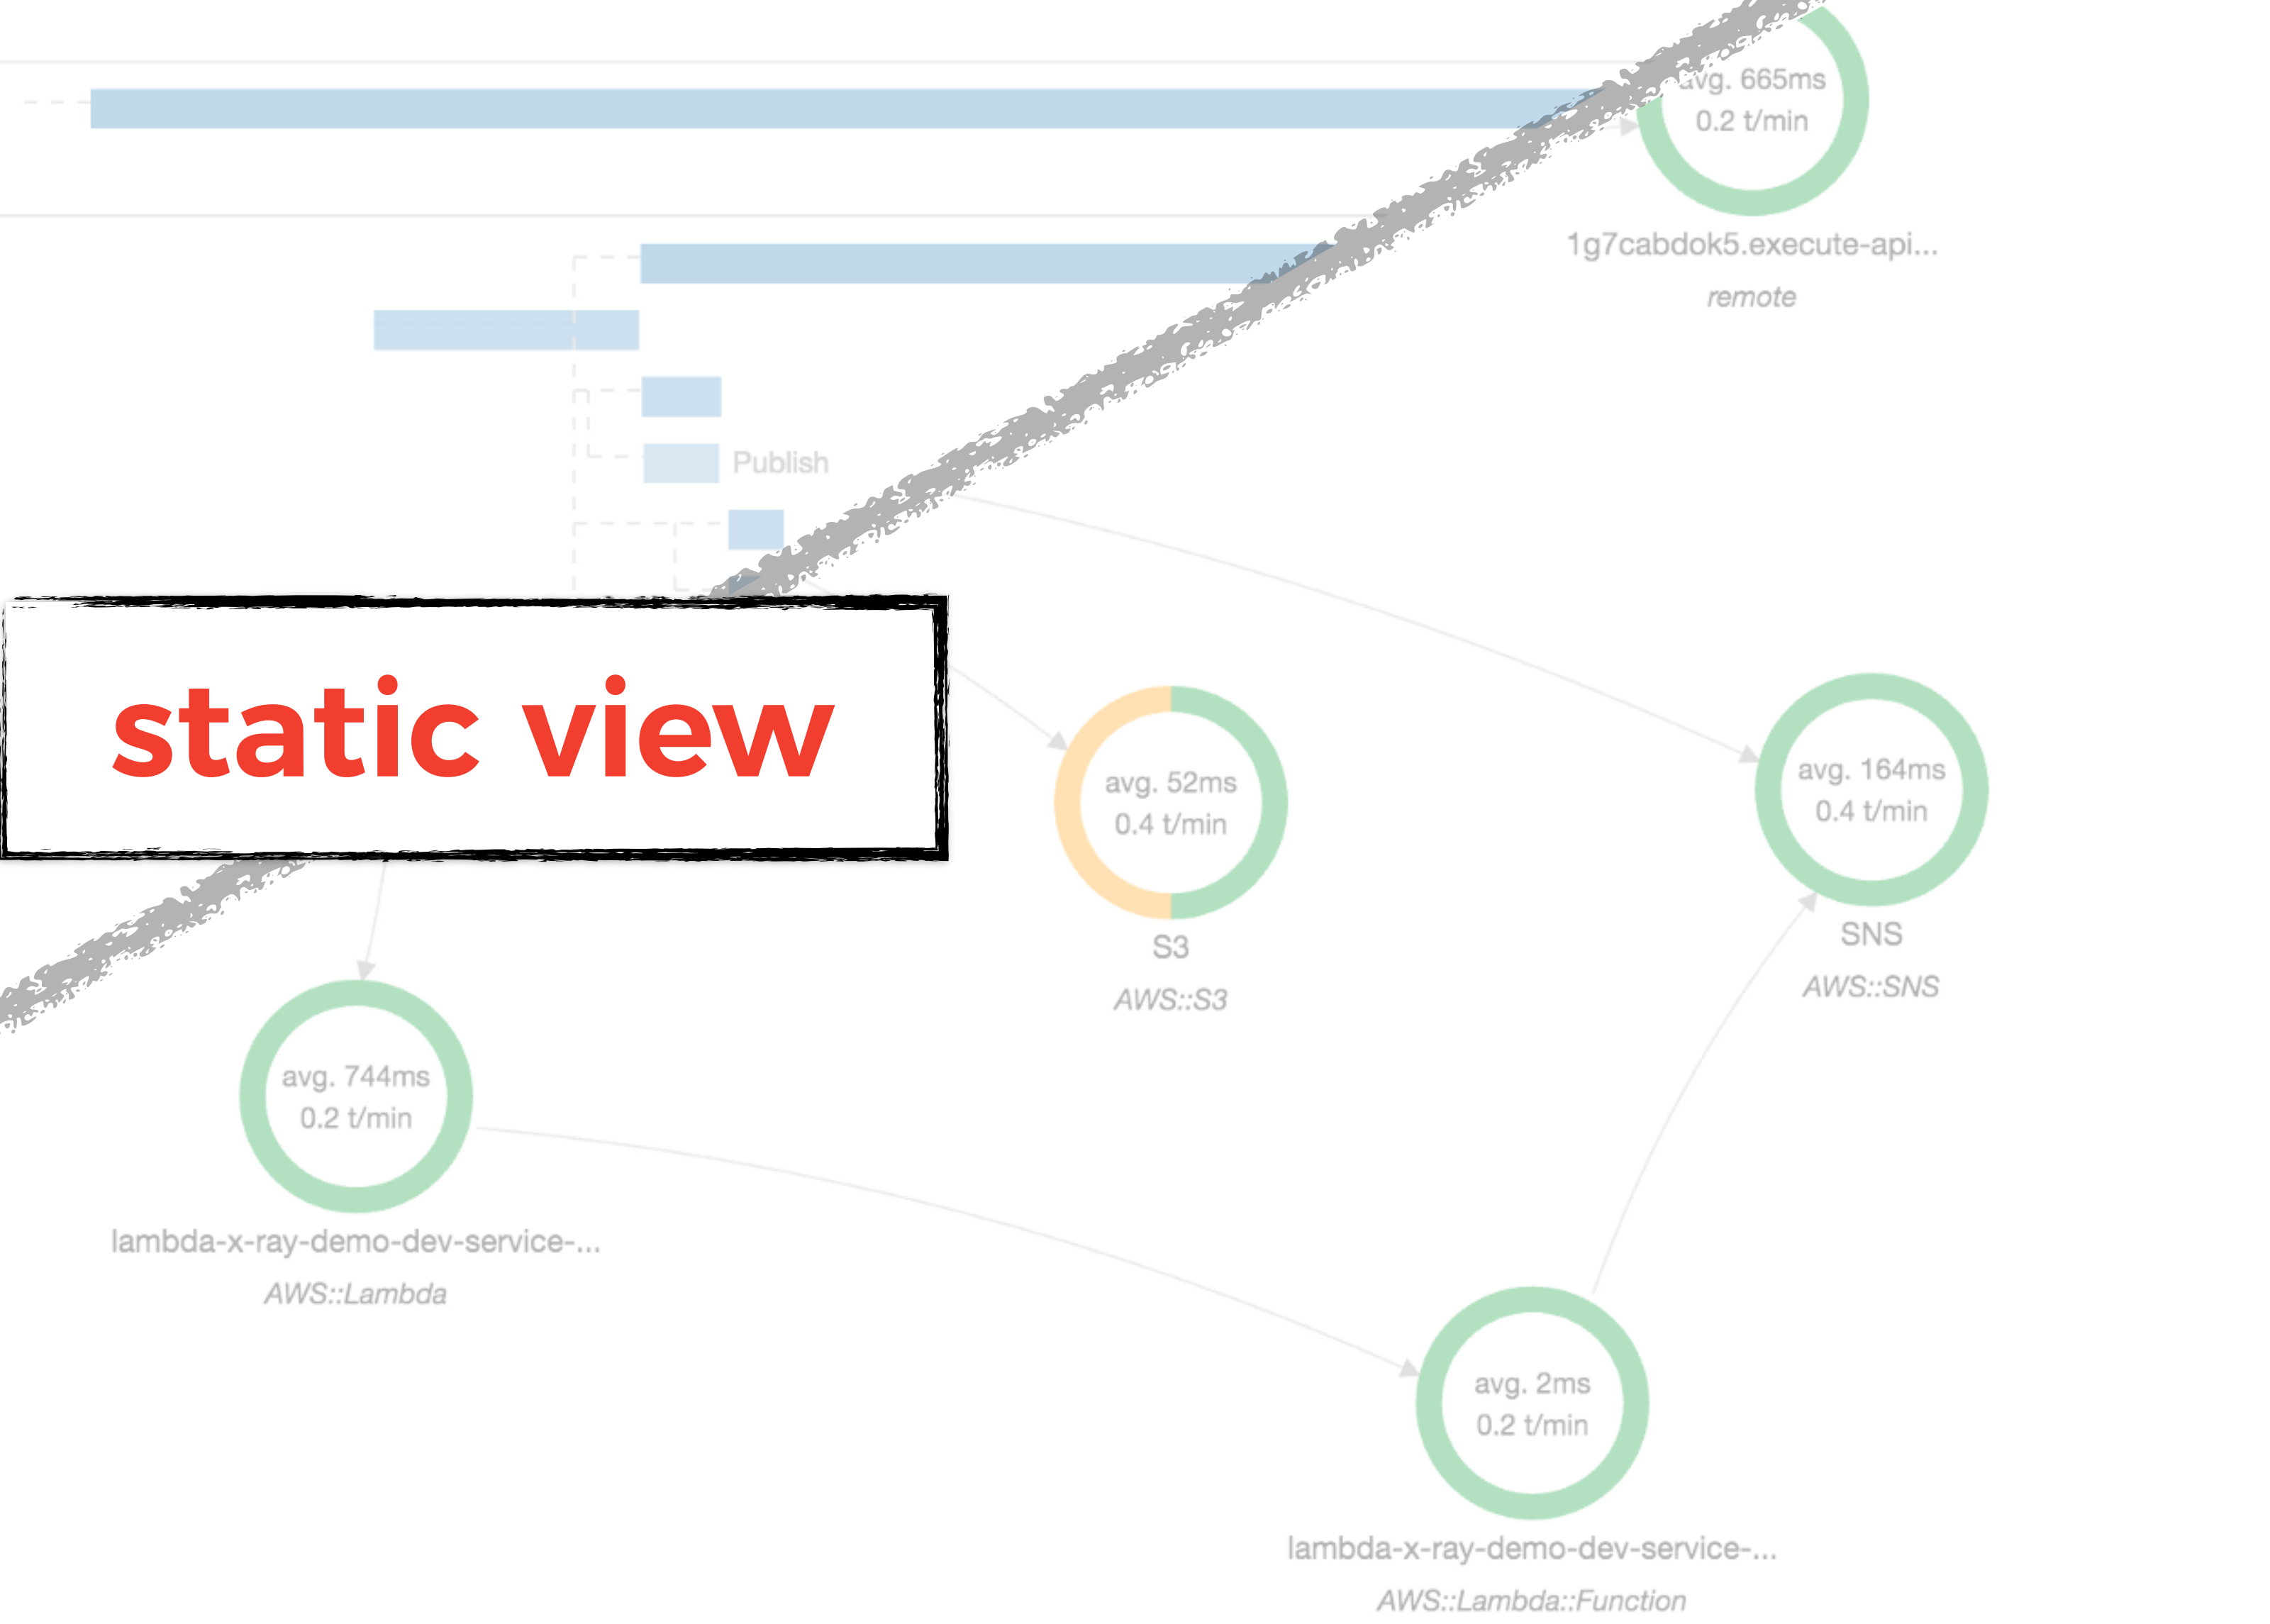
▼ lambda-x-ray-demo-dev-service-c AWS::Lambda			
lambda-x-ray-demo-dev-service-c			
▼ lambda-x-ray-demo-dev-service-c AWS::Lambda::Function			
lambda-x-ray-demo-dev-service-c			
Initialization			
## publishing			
SNS			



Name	Res.	Duration	Status	0.0ms	500ms	1.0s	1.5s	2.0s	2.5s	3.0s	3.5s	4.0s
------	------	----------	--------	-------	-------	------	------	------	------	------	------	------

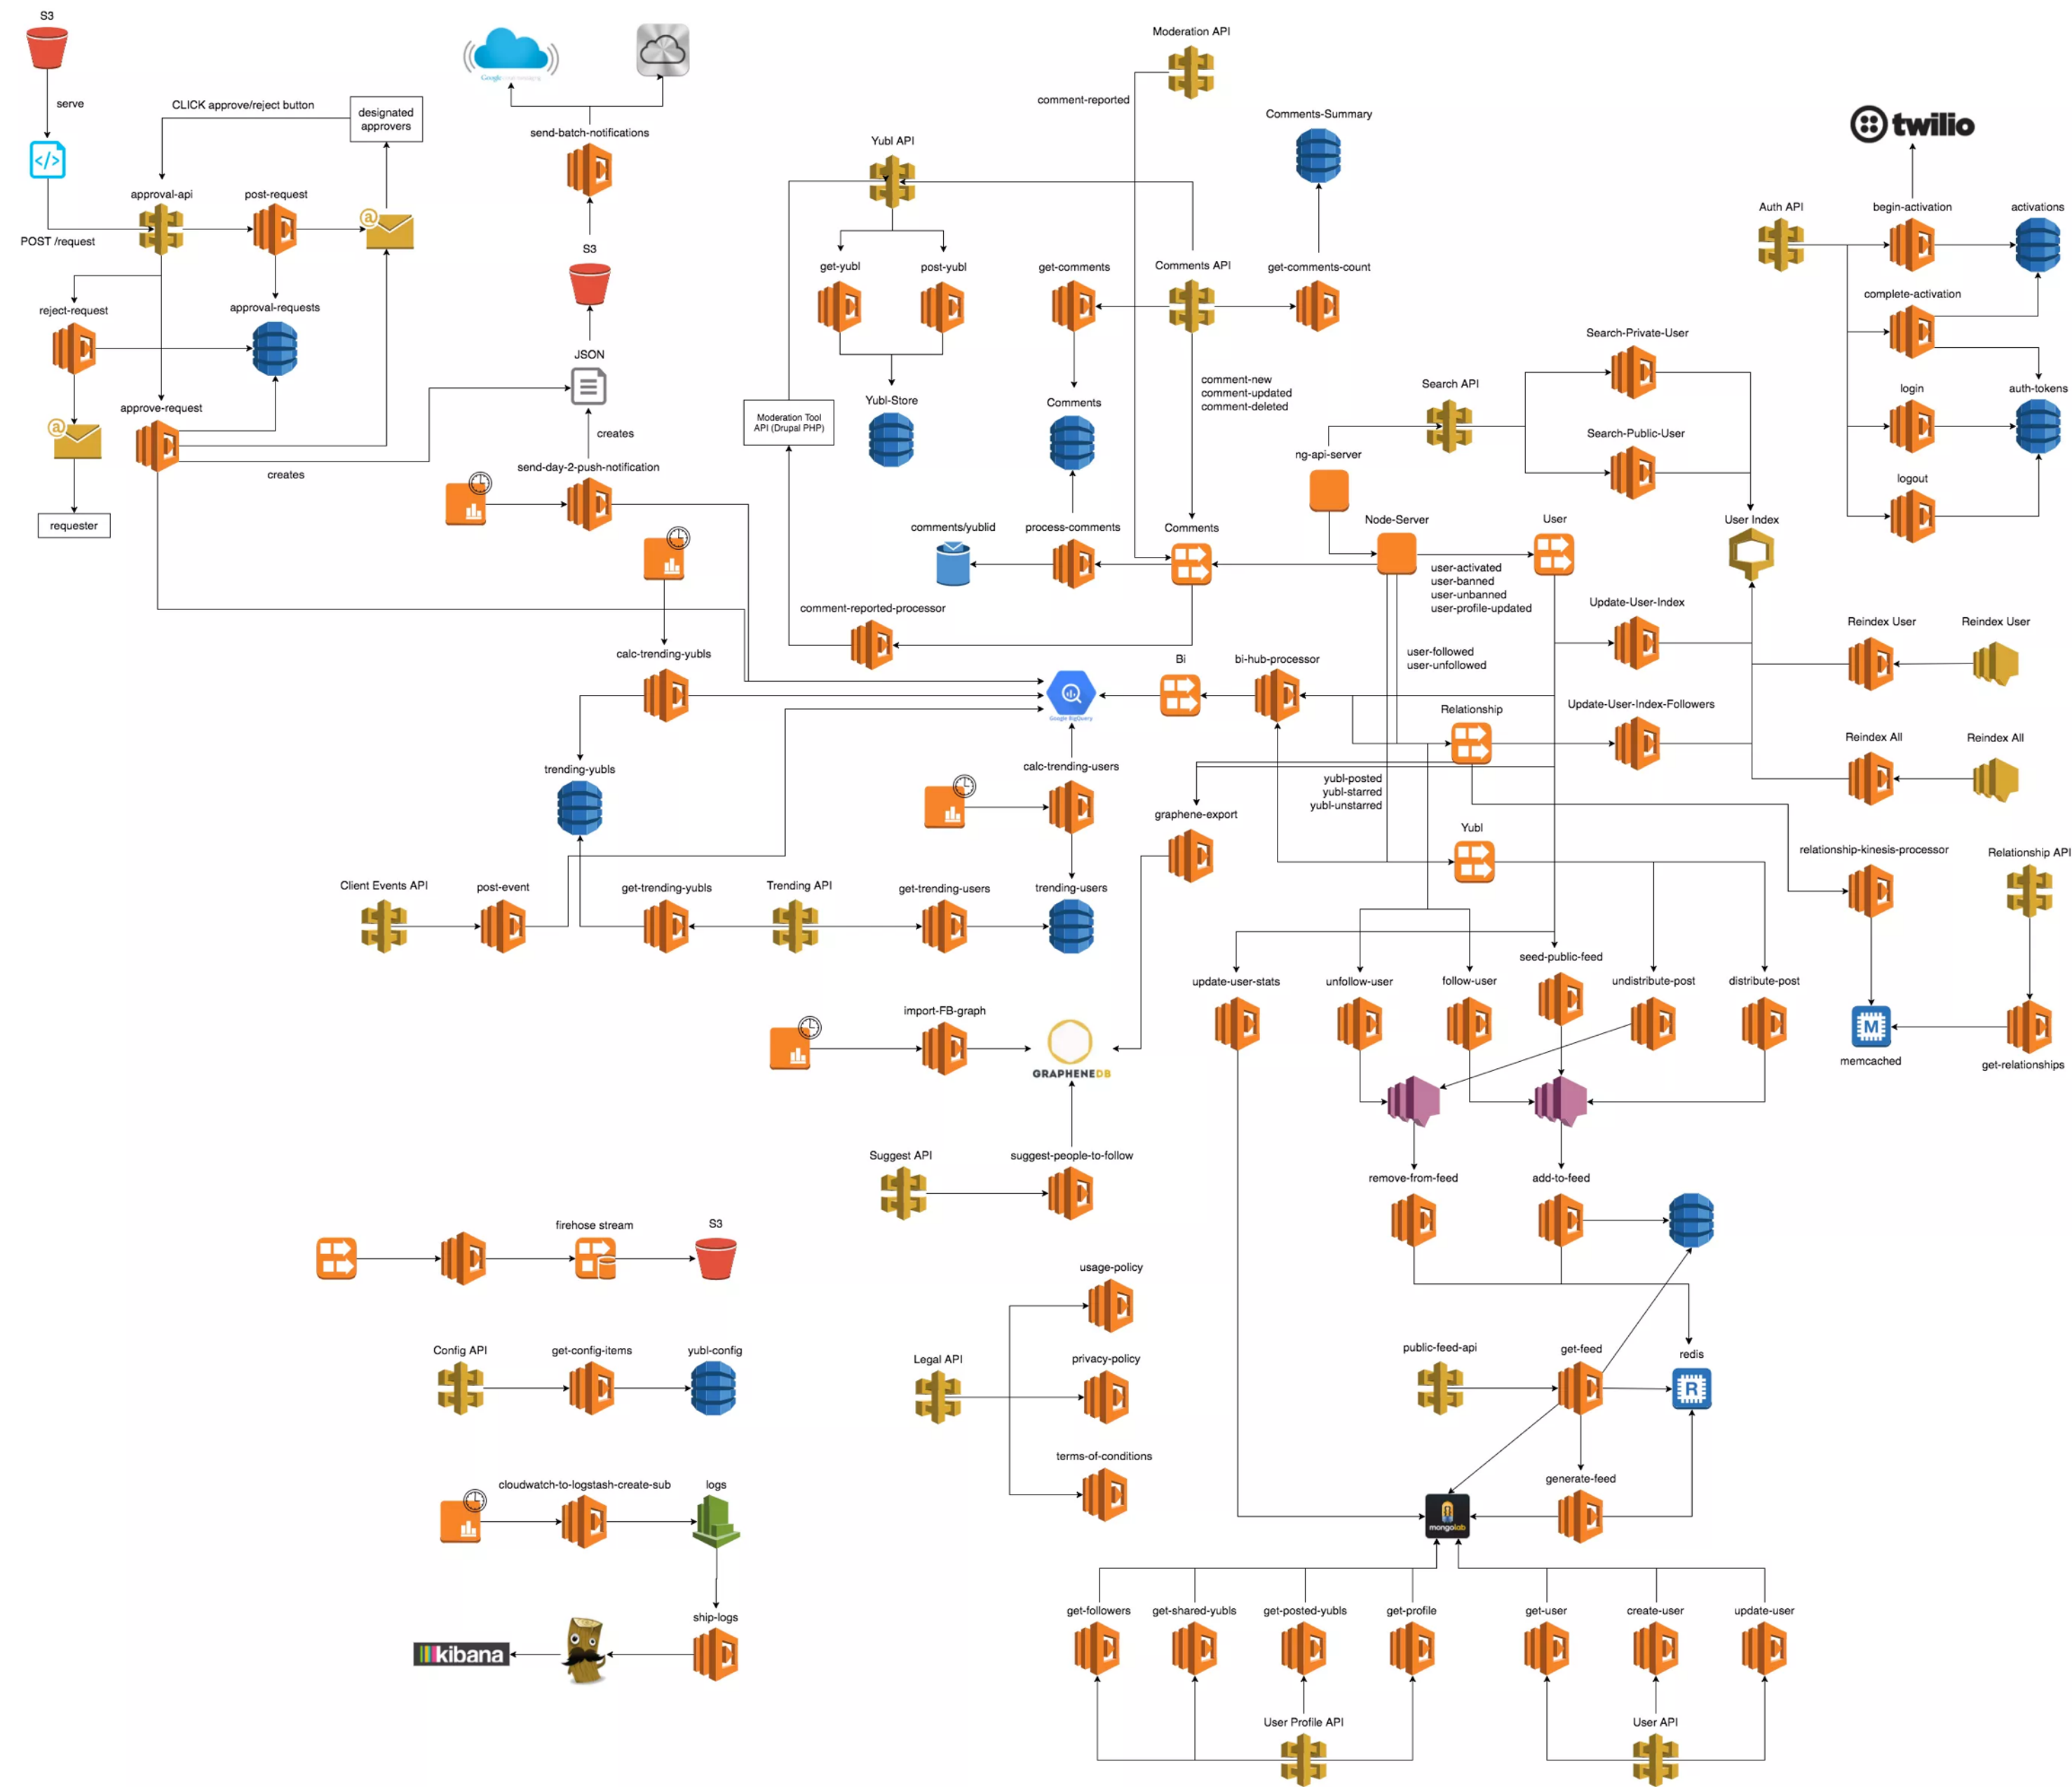
▼ lambda-x-ray-demo-dev-service-a AWS::Lambda			
lambda-x-ray-demo-dev-service-a	200	3.2 sec	✓
▼ lambda-x-ray-demo-dev-service-a AWS::Lambda::Function			
lambda-x-ray-demo-dev-service-a	-	2.1 sec	✓
Initialization	-	529 ms	✓
## publishing to SNS	-	158 ms	✓
SNS	200	151 ms	✓
## accessing S3	-	110 ms	✓
S3	404	65.0 ms	!
S3	200	40.0 ms	✓
## accessing DynamoDB	-	72.0 ms	✓
DynamoDB	200	36.0 ms	✓
DynamoDB	200	32.0 ms	✓
## invoking service-c	-	1.2 sec	✓
Lambda	200	1.2 sec	✓
## calling service b	-	550 ms	✓
1g7cabdok5.execute-api.us-east-1.amazonav	200	548 ms	✓

static view

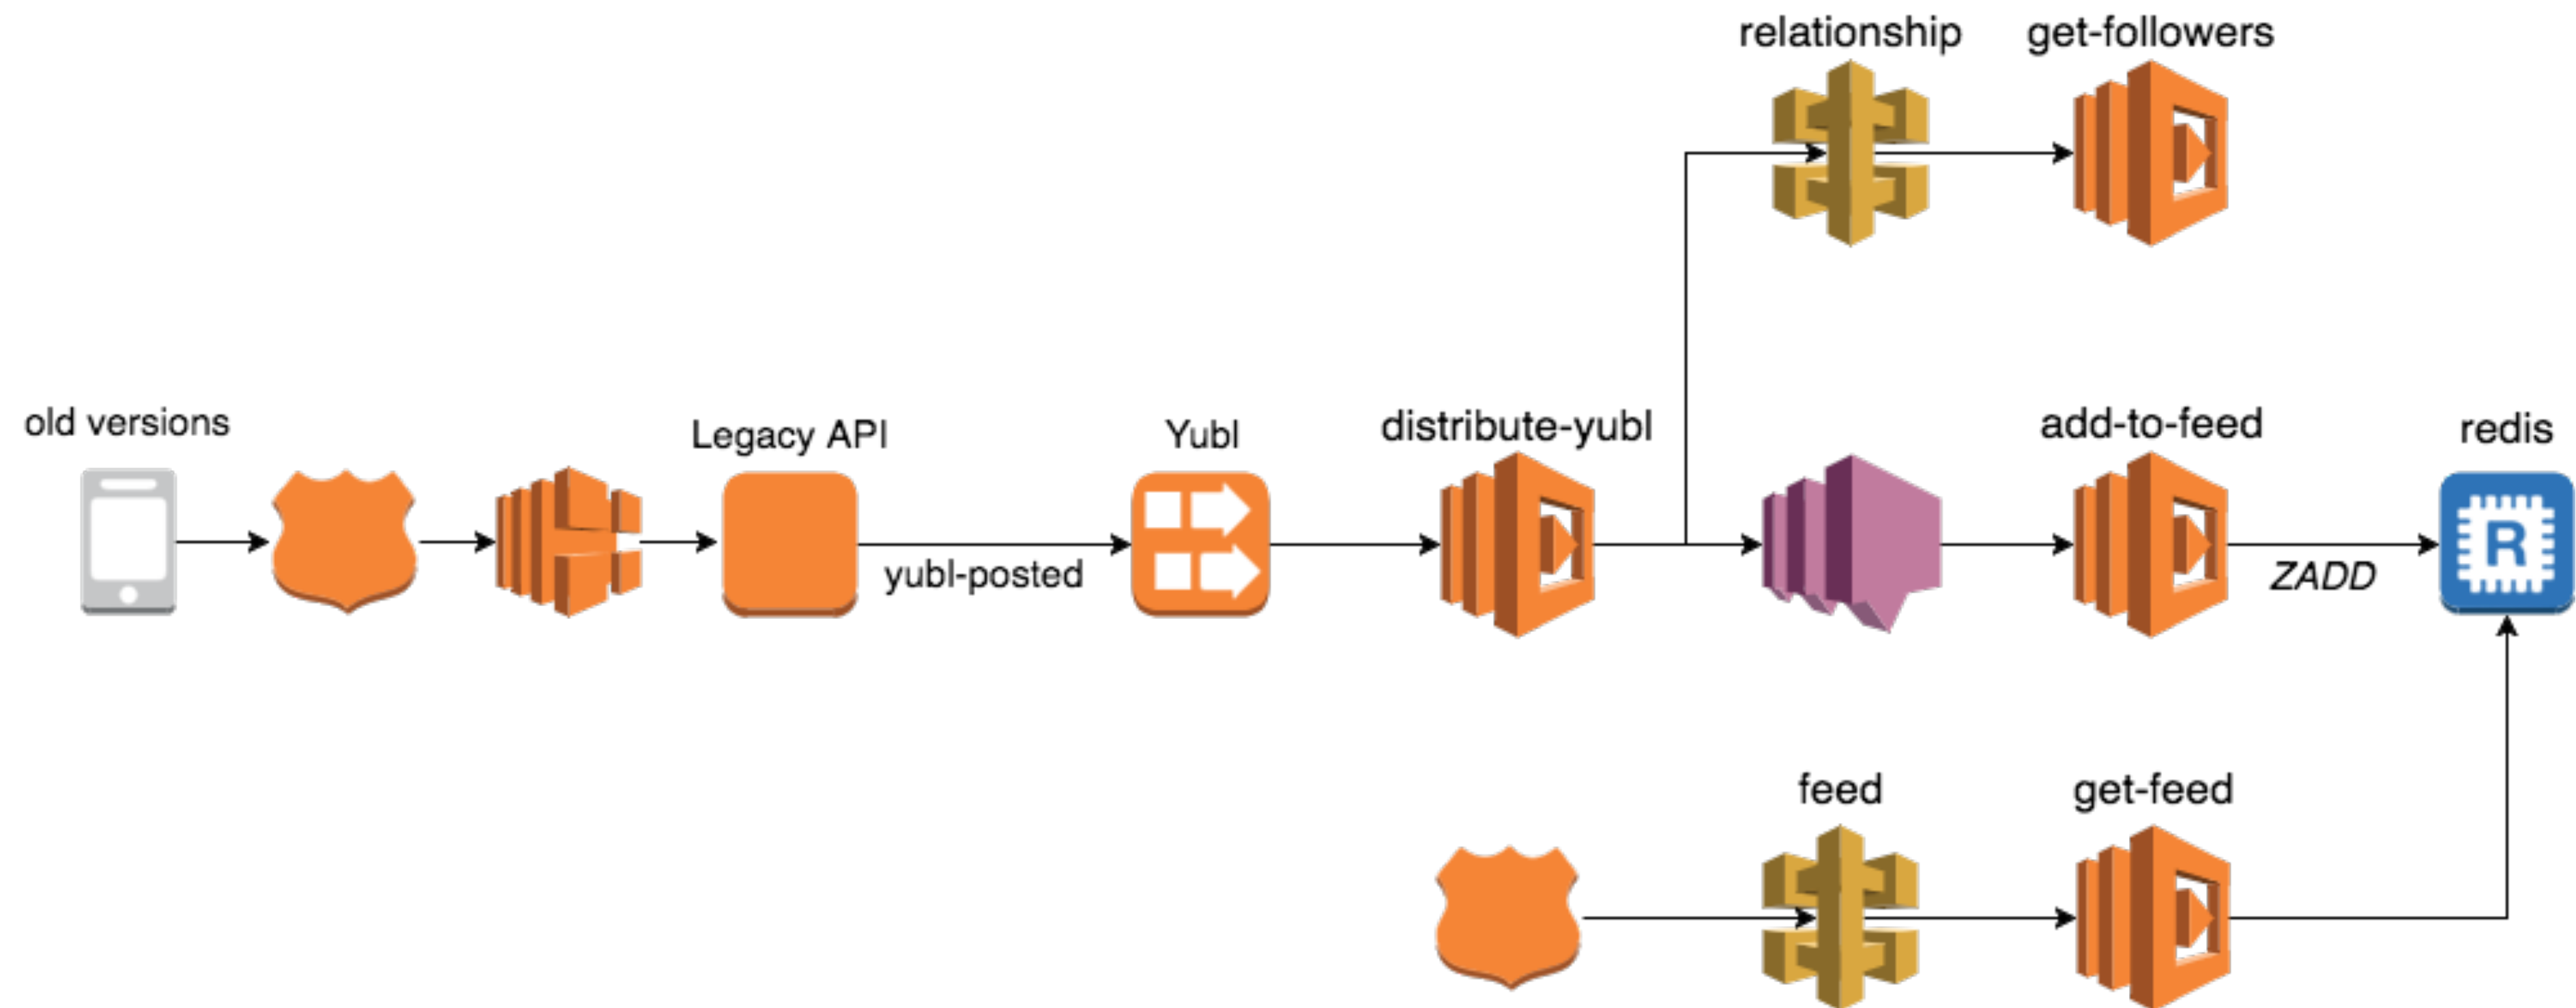


▼ lambda-x-ray-demo-dev-service-c AWS::Lambda			
lambda-x-ray-demo-dev-service-c			
▼ lambda-x-ray-demo-dev-service-c AWS::Lambda::Function			
lambda-x-ray-demo-dev-service-c			
Initialization			
## publishing to			
SNS			

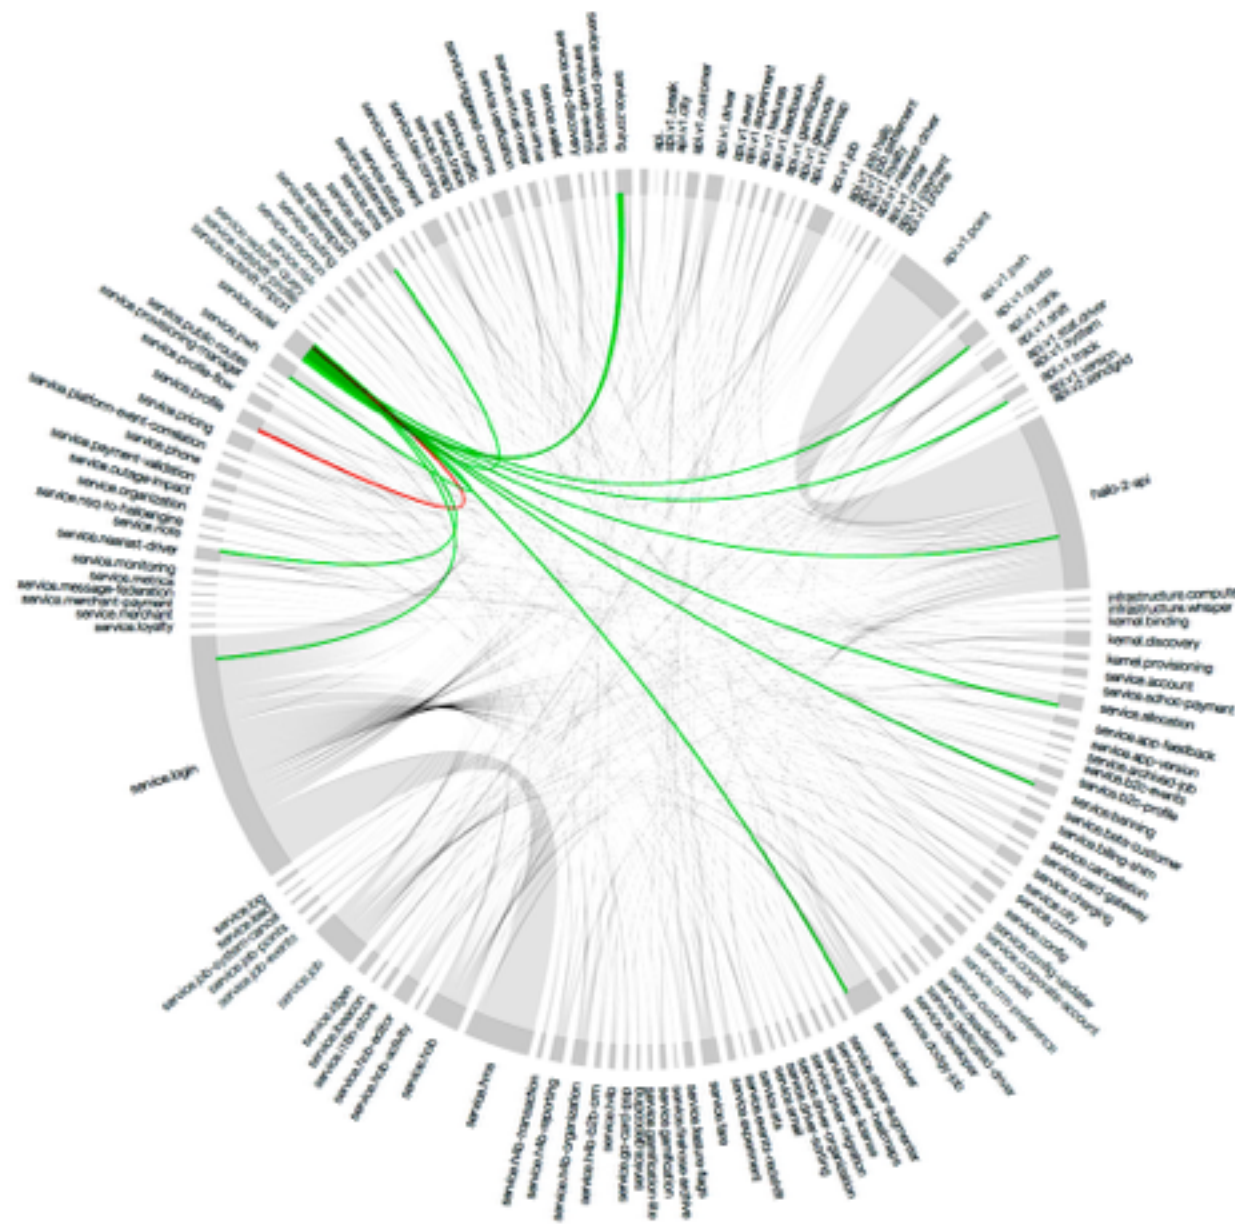
our tools need to do more to help us with understanding & debugging our distributed system, not just what happens inside one function



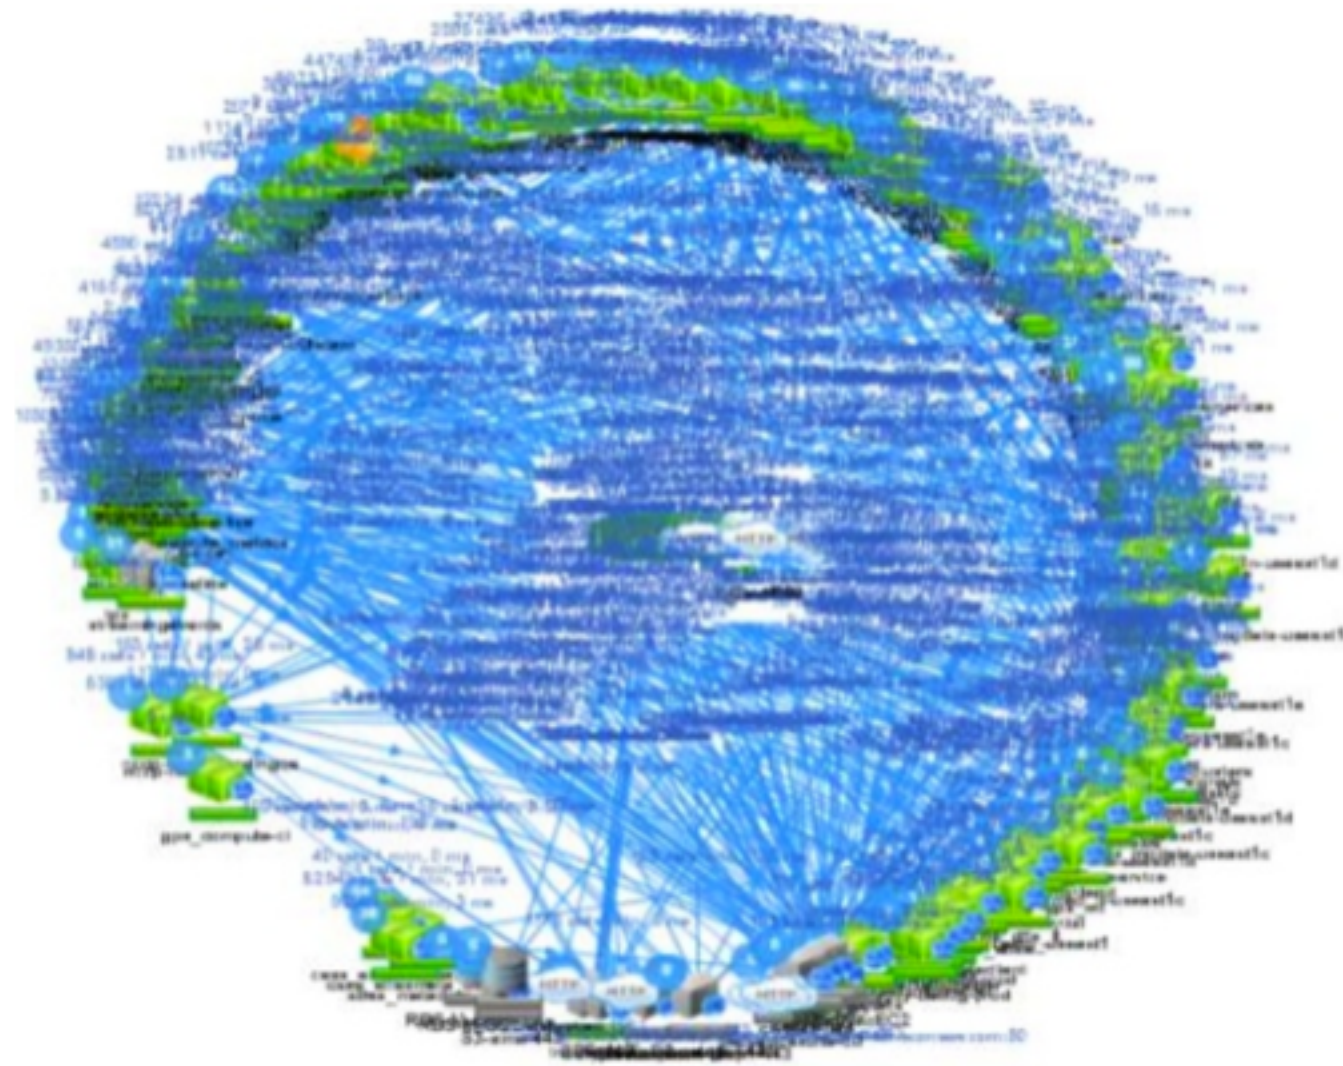
“one user action/vertical slice through the system”



450+ microservices

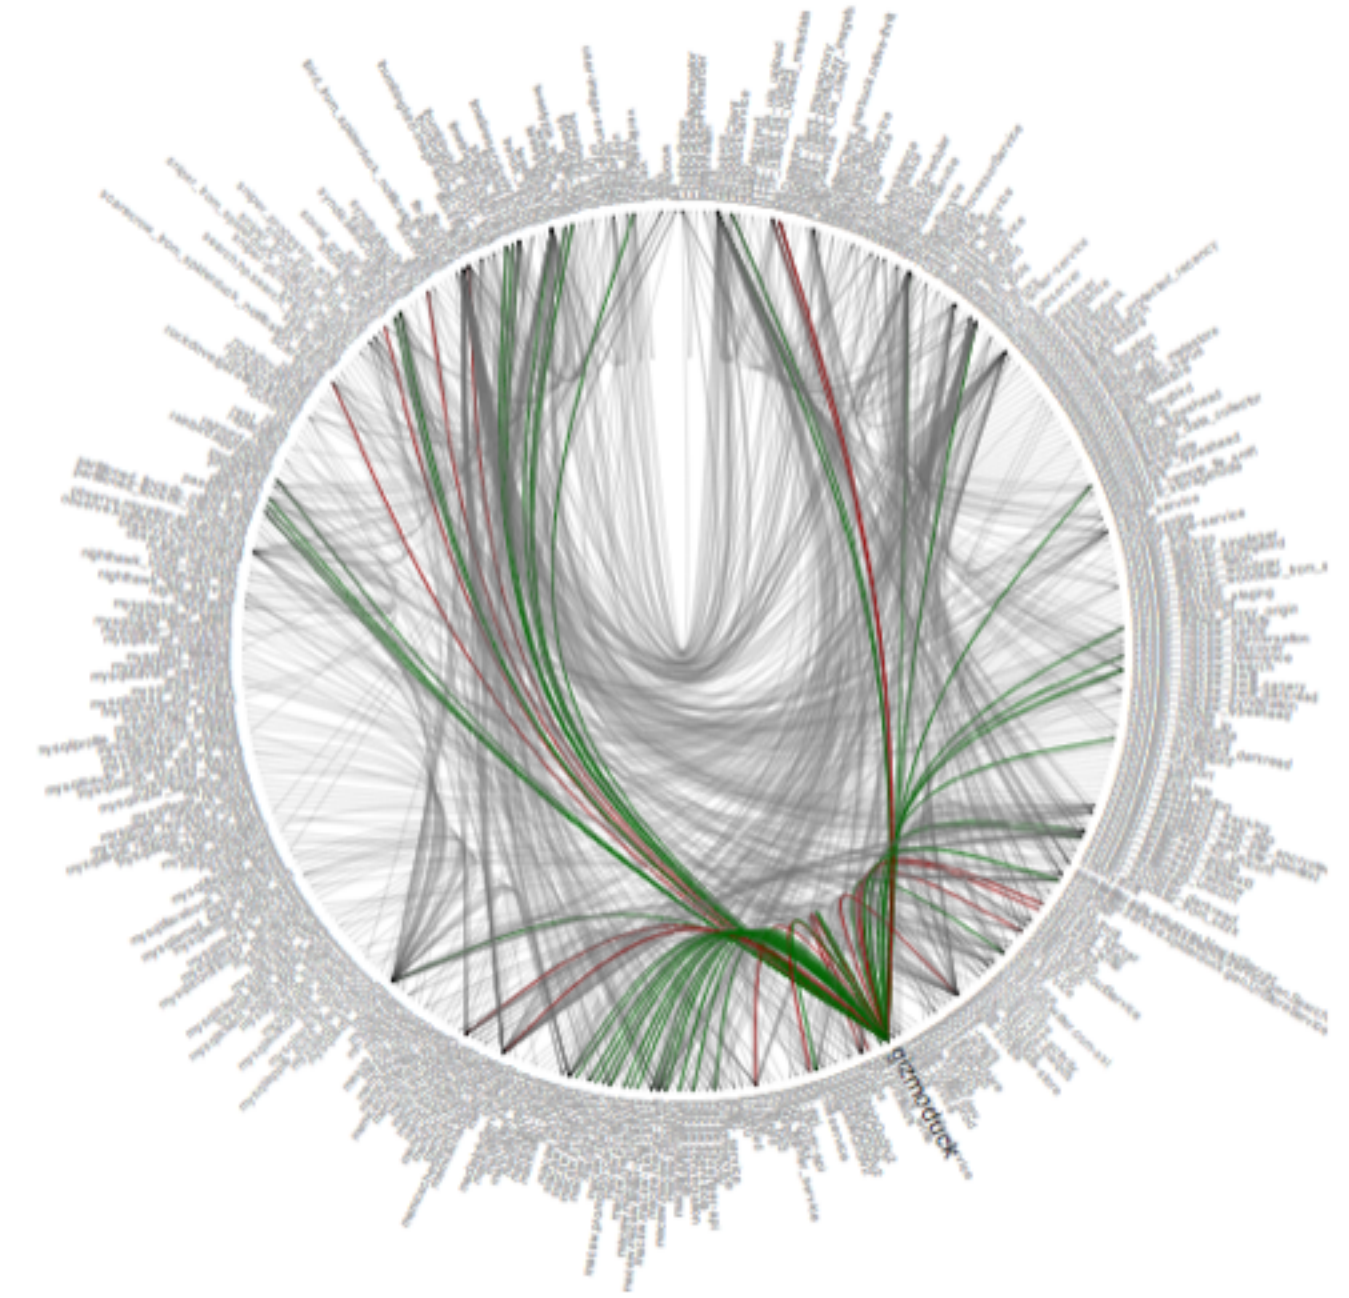


500+ microservices



NETFLIX

500+ microservices

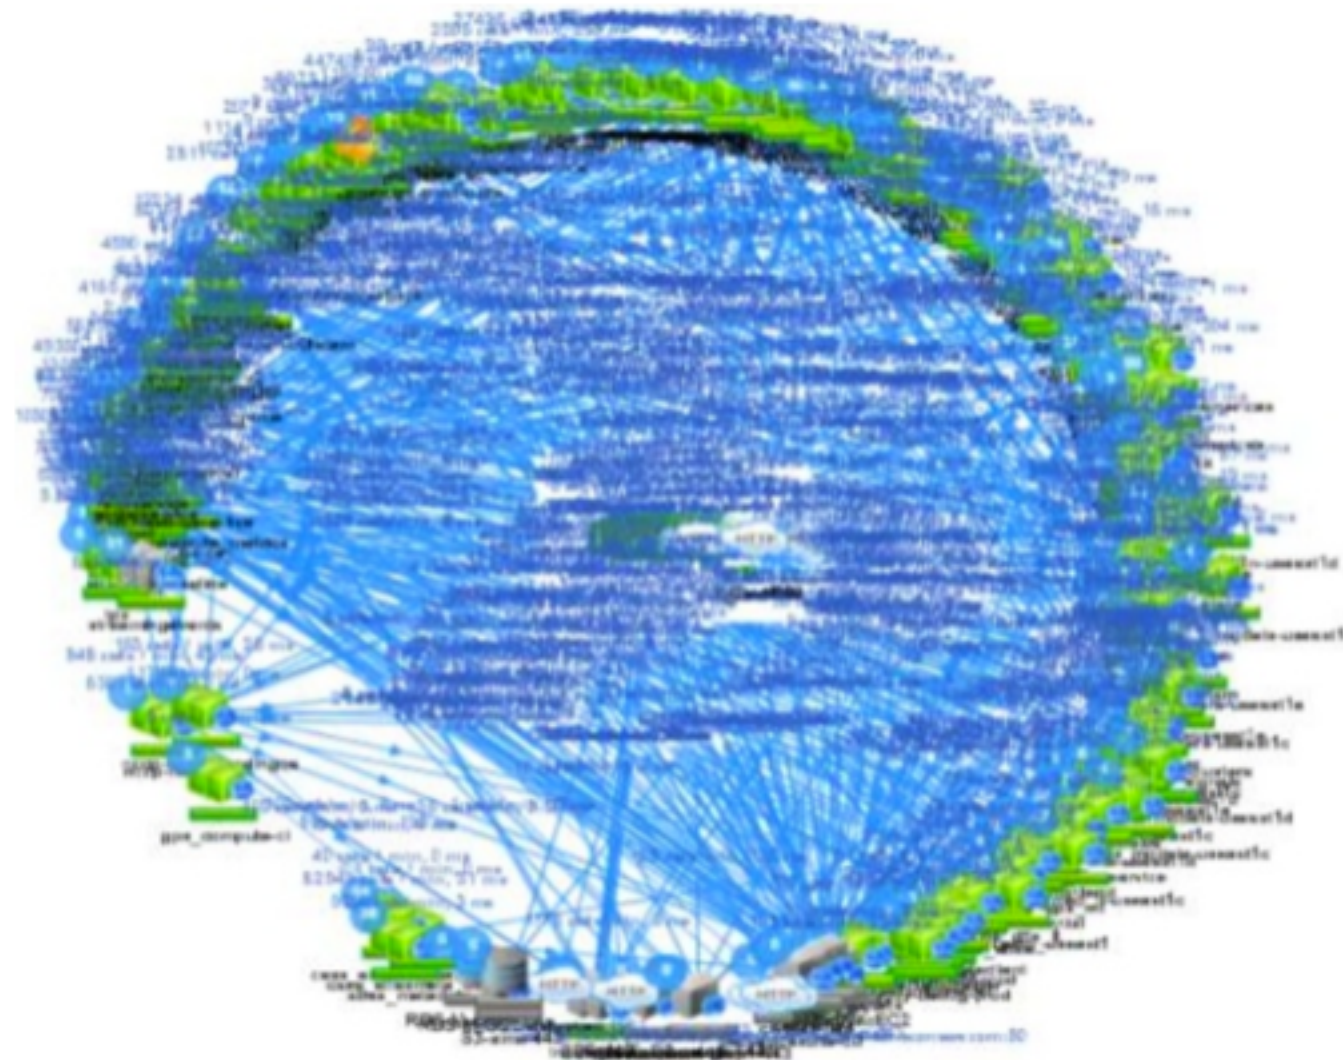
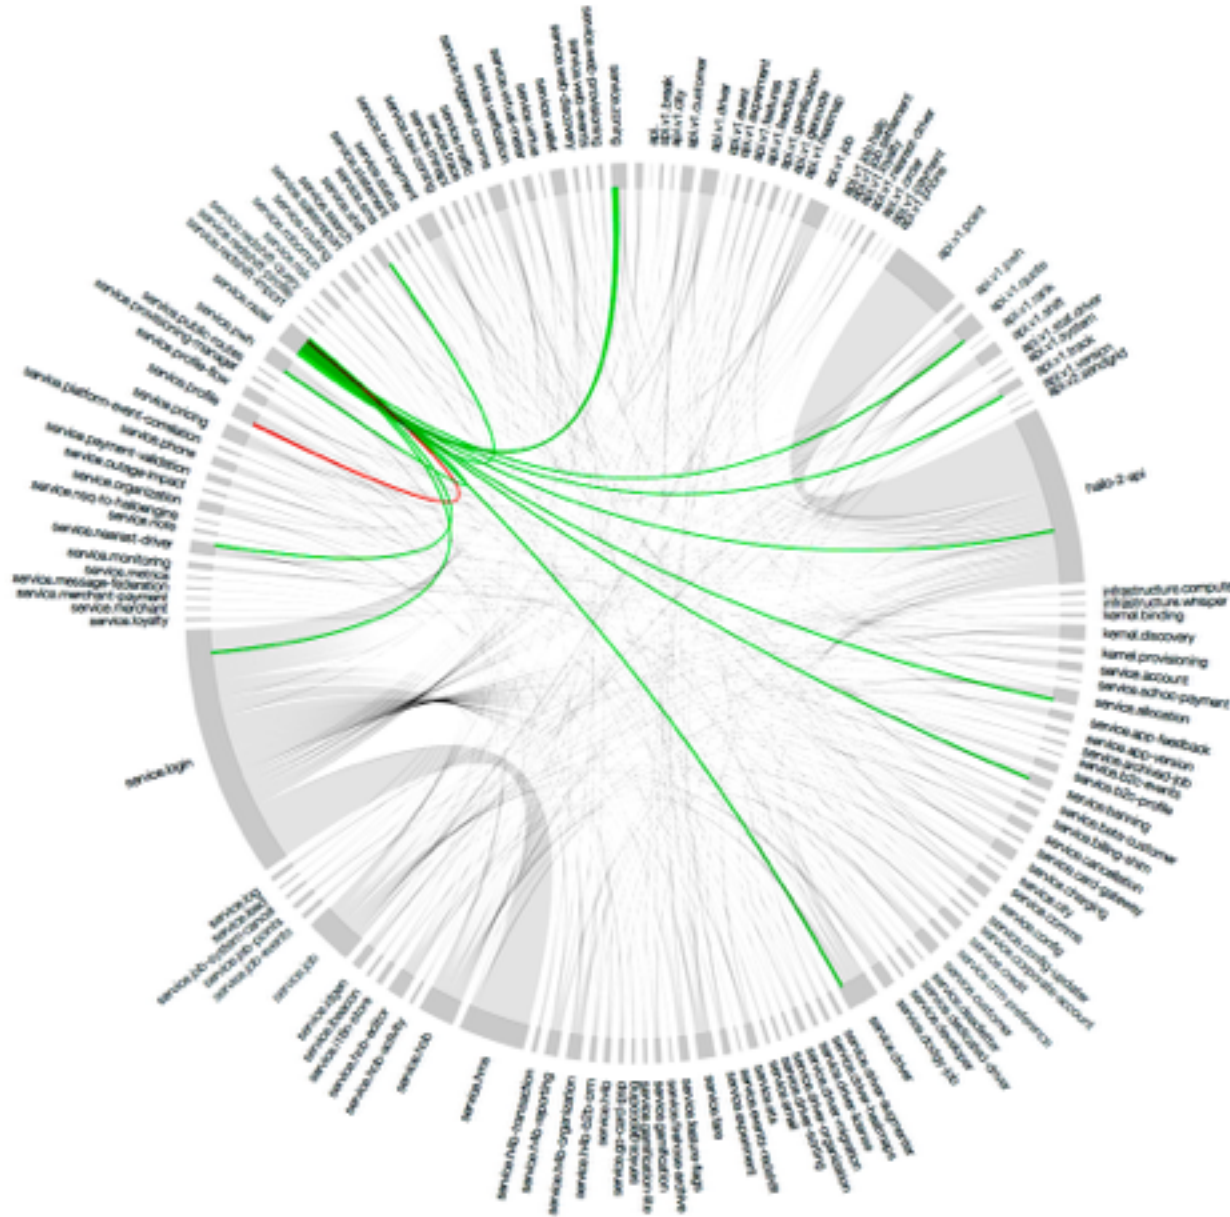


microservices death stars circa 2015

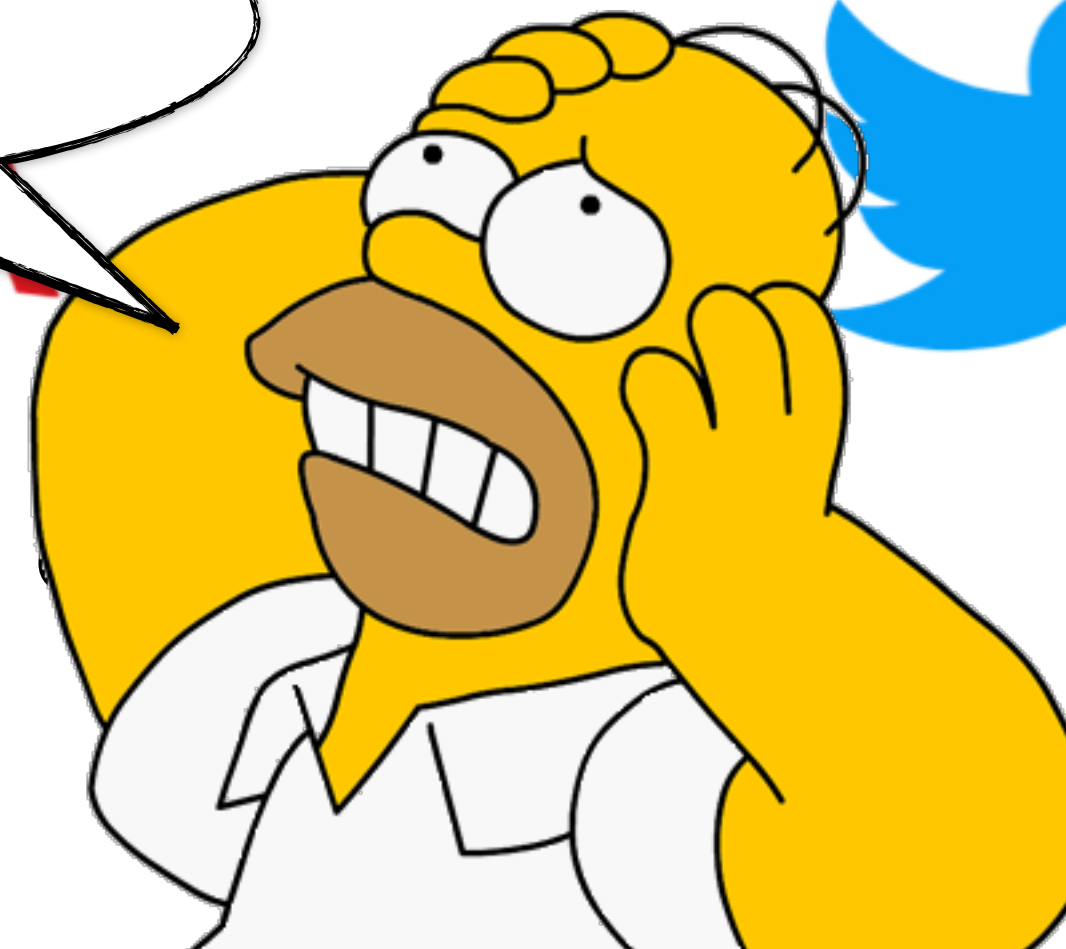
450+ microservices

500+ microservices

500+ microservices



microservices death stars

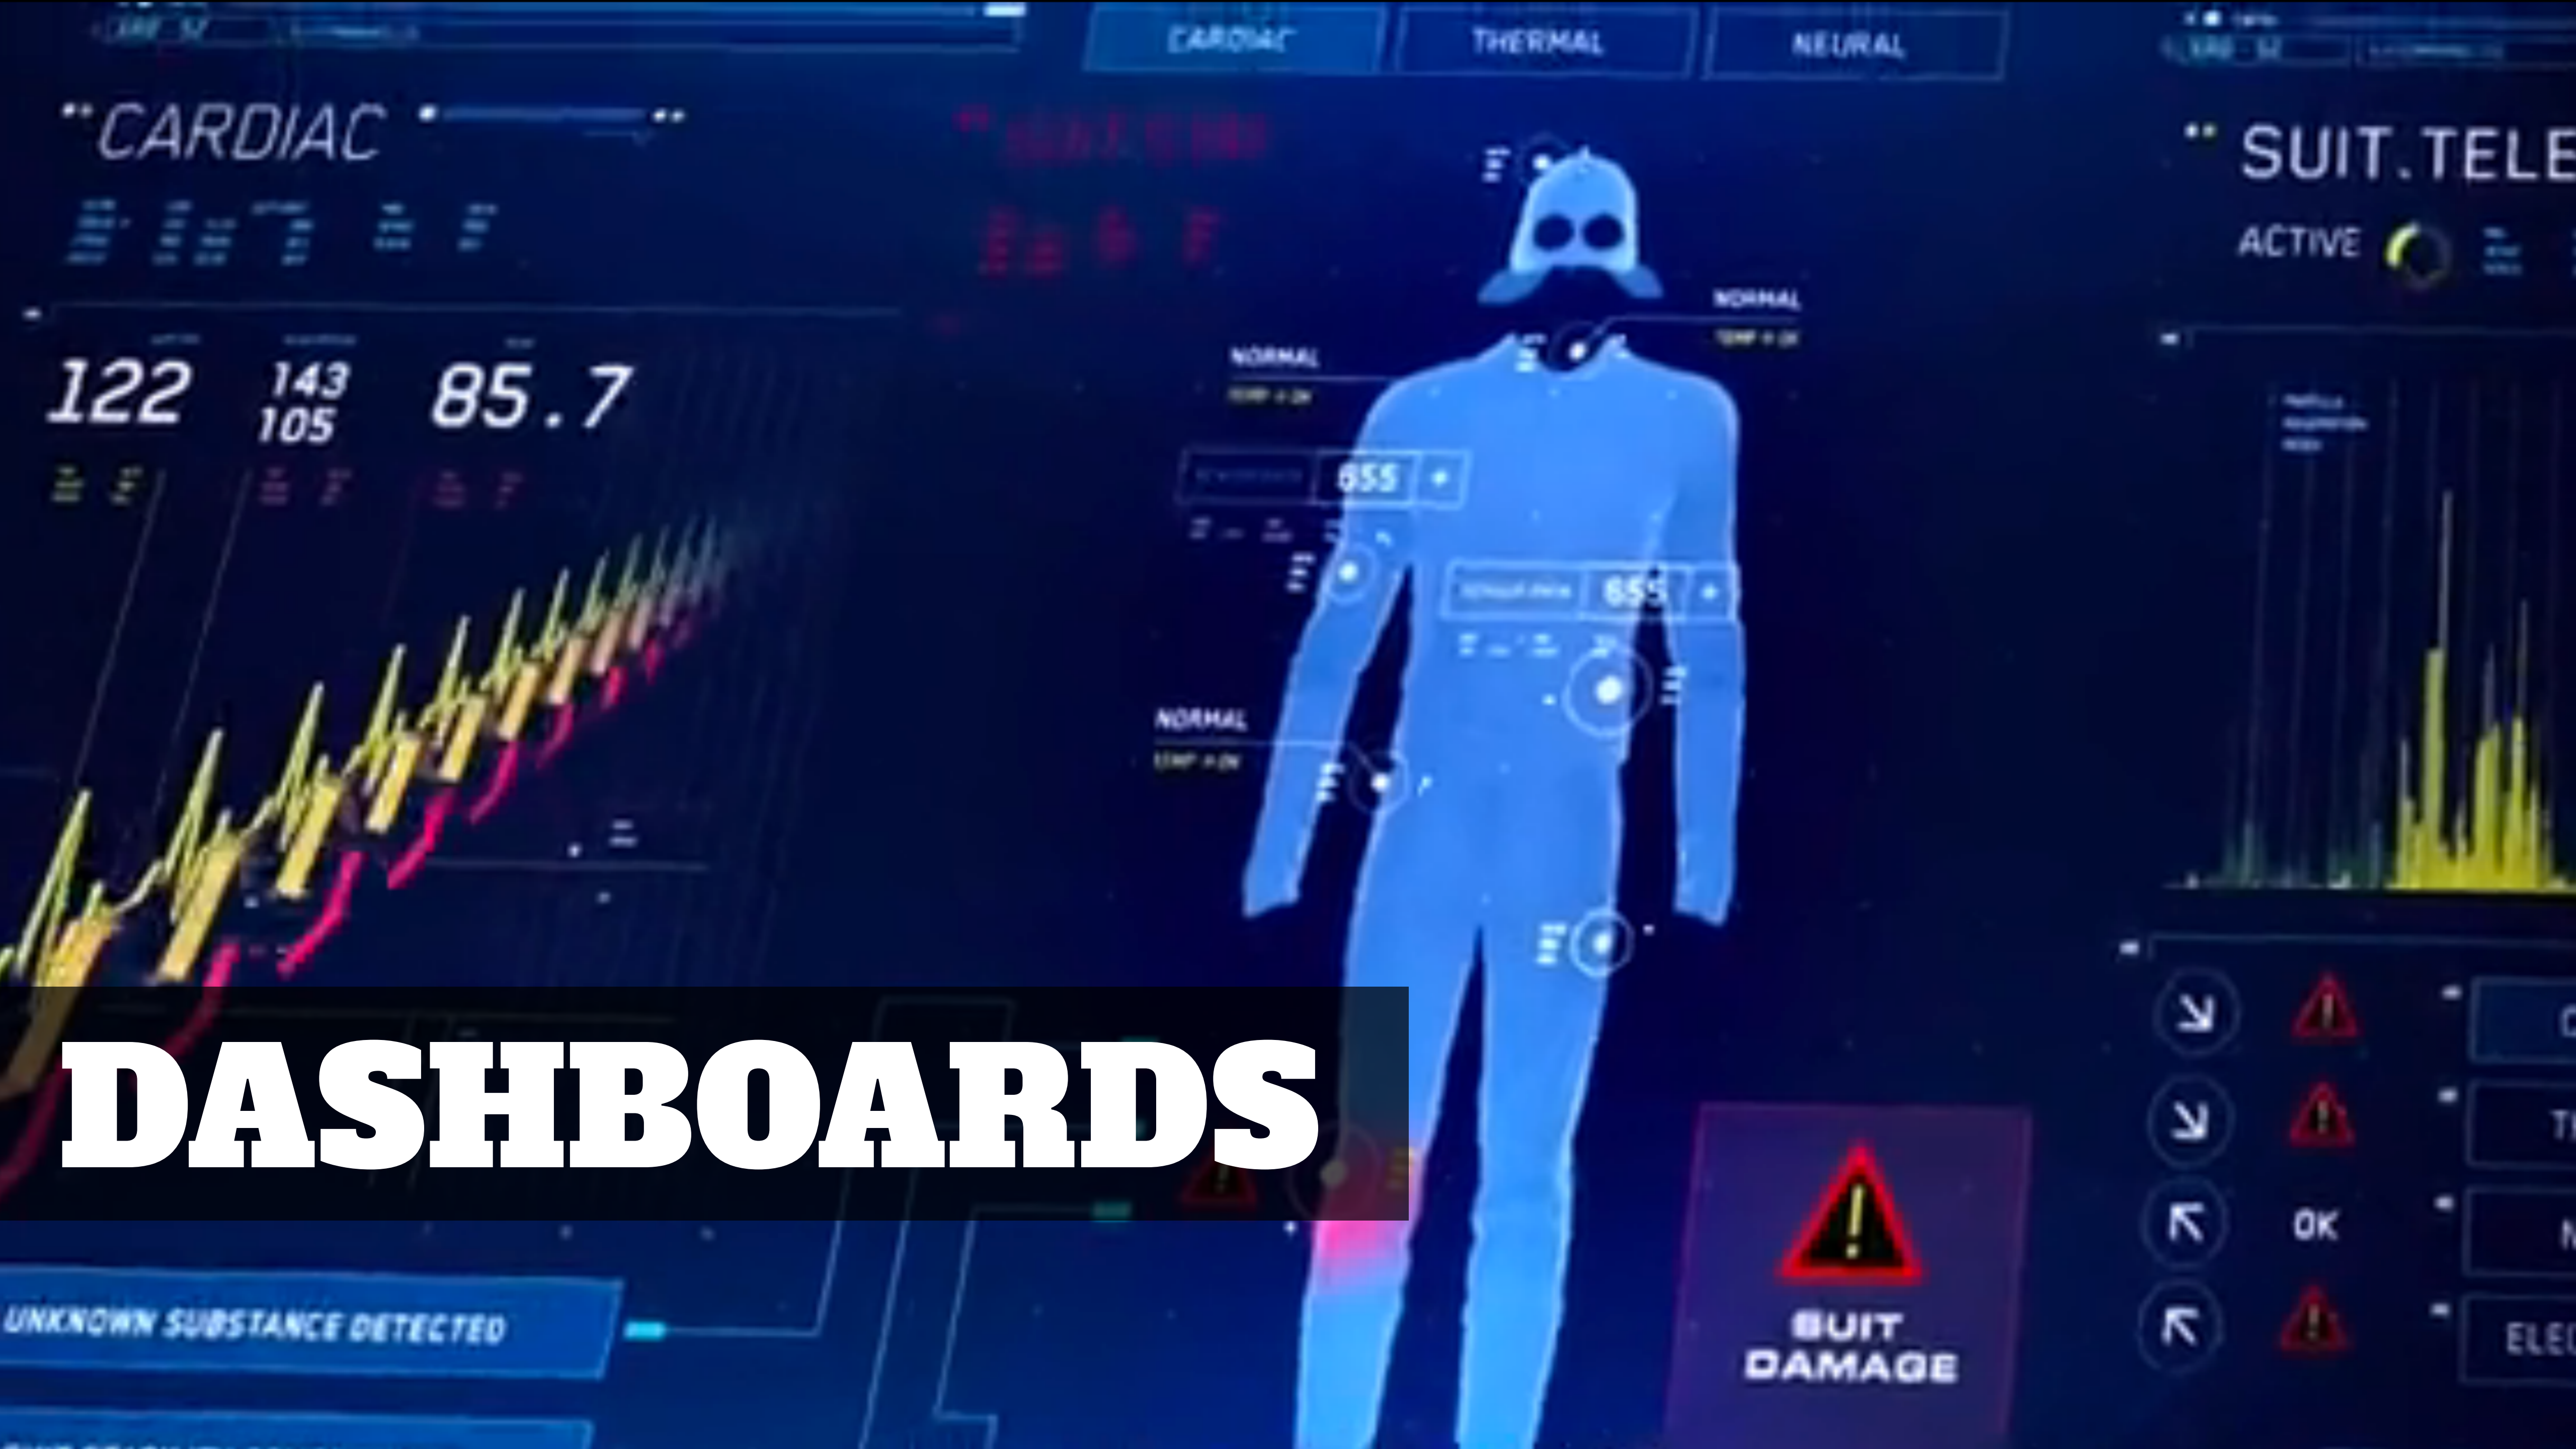




THE FUTURE IS NOW



WARNING: this is part fiction, part inspired by new tools



CARDIAC

THERMAL

NEURAL

CARDIAC

SUIT.TELE

ACTIVE

122 143 85.7
105

NORMAL

NORMAL

NORMAL

655

655

DASHBOARDS



SUIT DAMAGE

UNKNOWN SUBSTANCE DETECTED

Control panel with directional arrows (up, down, left, right), a red warning triangle, and the text 'OK'. Below these are buttons labeled 'C', 'T', 'N', and 'ELEC'.

- Dashboard
- Reports
- Team
- Configure
- History

All teams - quarter-to-quarter

+14%
(QTD)

Monthly Trend
(last 3 months)

Last updated 4 days ago

Product team - year-over-year

-6%
(QTD)

Yearly Trend
(last 3 years)

Last updated 3 days ago

Top and bottom 3 teams - percentage of goals (YTD)

Front-end	91%	Product	65%
UX design	84%	Big Data	53%
QA	72%	Analytics	47%

Last updated 4 days ago

Quarterly team performance - last 2 years (percentage of goal)

PRODUCT | UX DESIGN | ENGINEERING

Monthly

Year	Q1	Q2	Q3	Q4	Q1	Q2	Q3	QTD
2015	43	31	29	68				
2016	78	47	63	54				

Last updated 2 days ago

- See more detail
- Add to report
- Edit card
- Replace

Meetings attended by team (QTD)

PRODUCT | UX DESIGN | ENGINEERING

73%
Attended

Last updated 1 day ago

Engineering performance (QTD)

NEW YORK | PUNE | SYDNEY

85%
On target

Last updated 2 days ago

\$monitor_name * ?



Overview

Success rate 1d

(No data)

Average size 2d

1.42 MiB

Load time 1d

2.49 s

Availability

Last 30 min... 30m

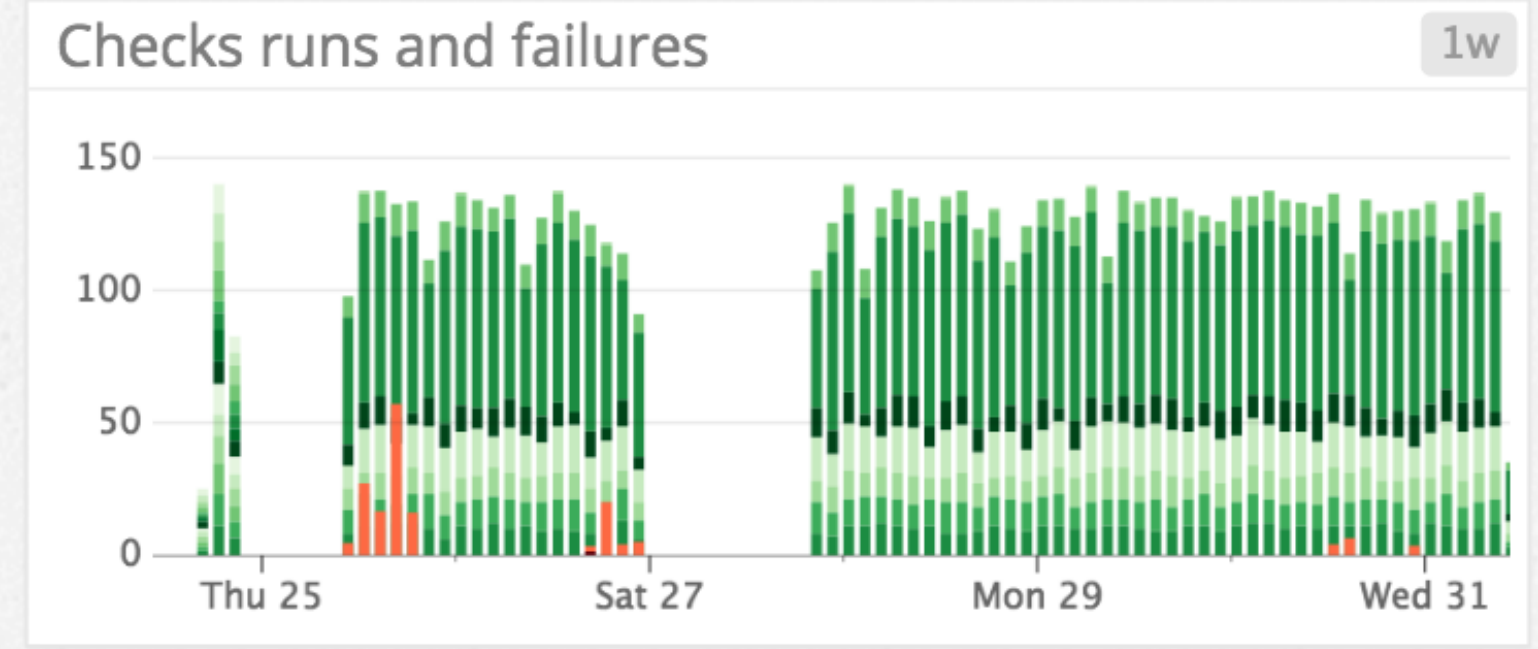
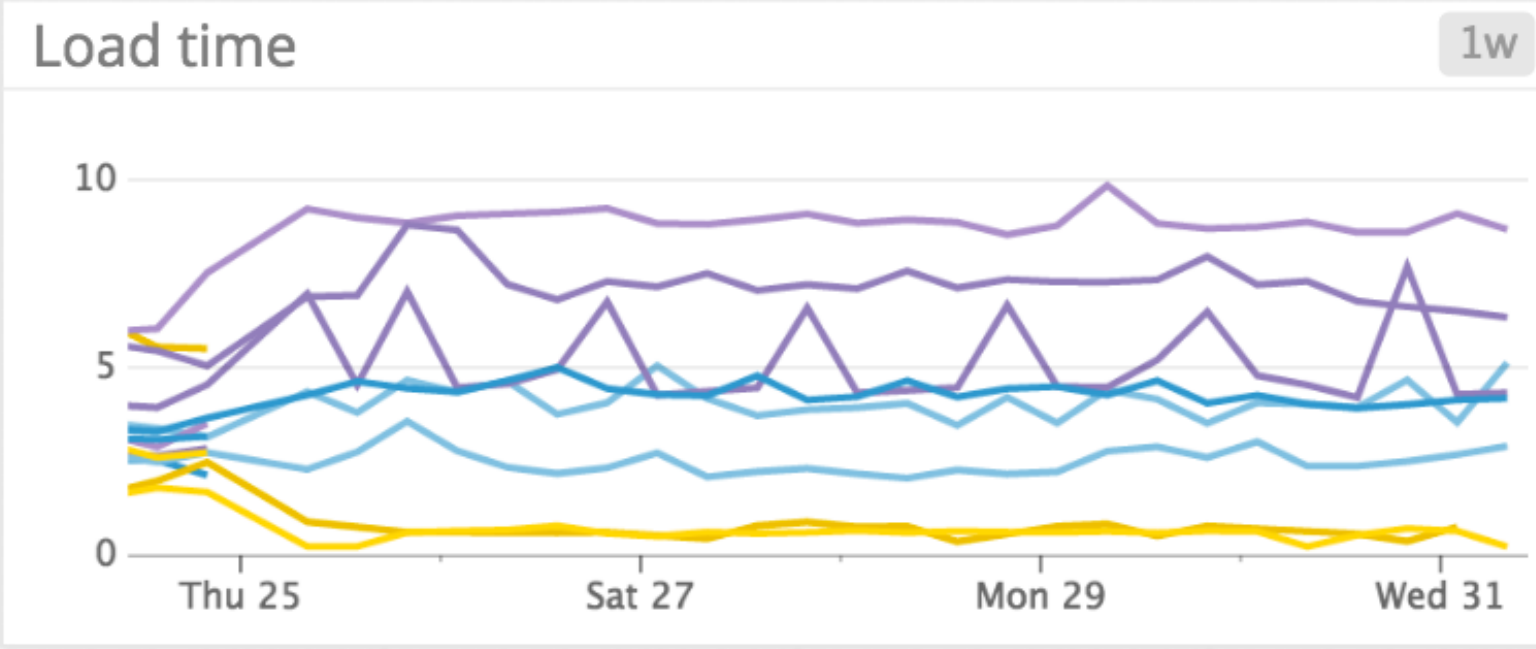
100%

Last 7 days 1w

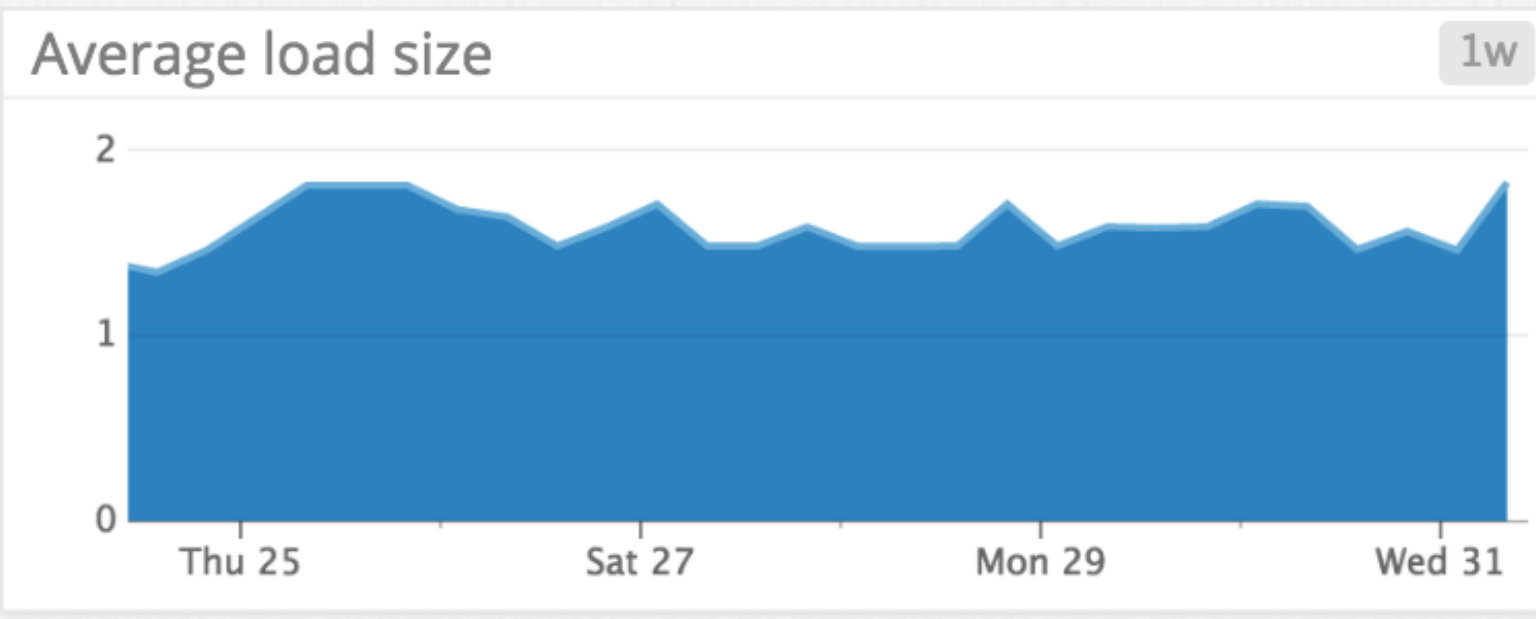
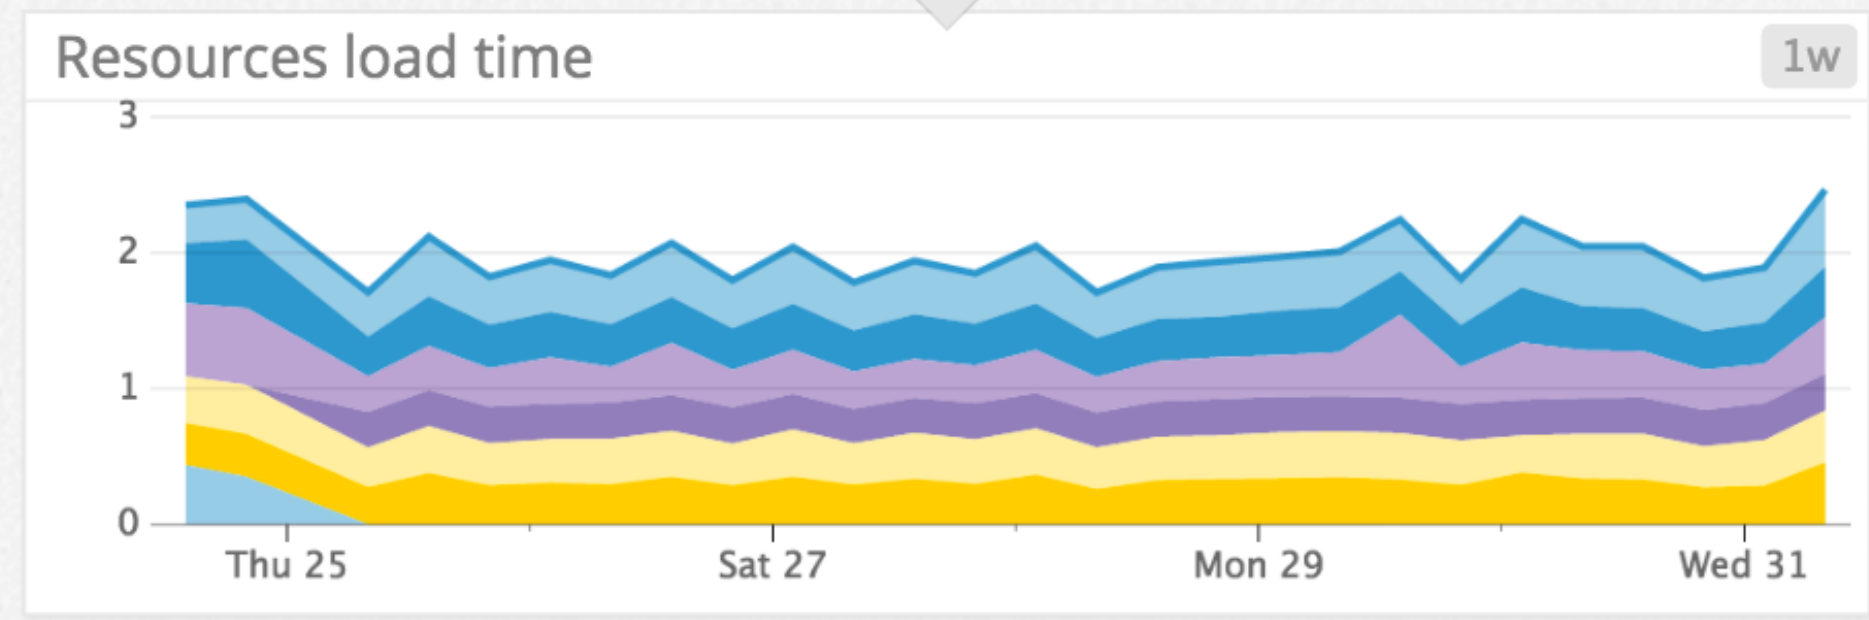
100%

Last 30 days 1mo

100%

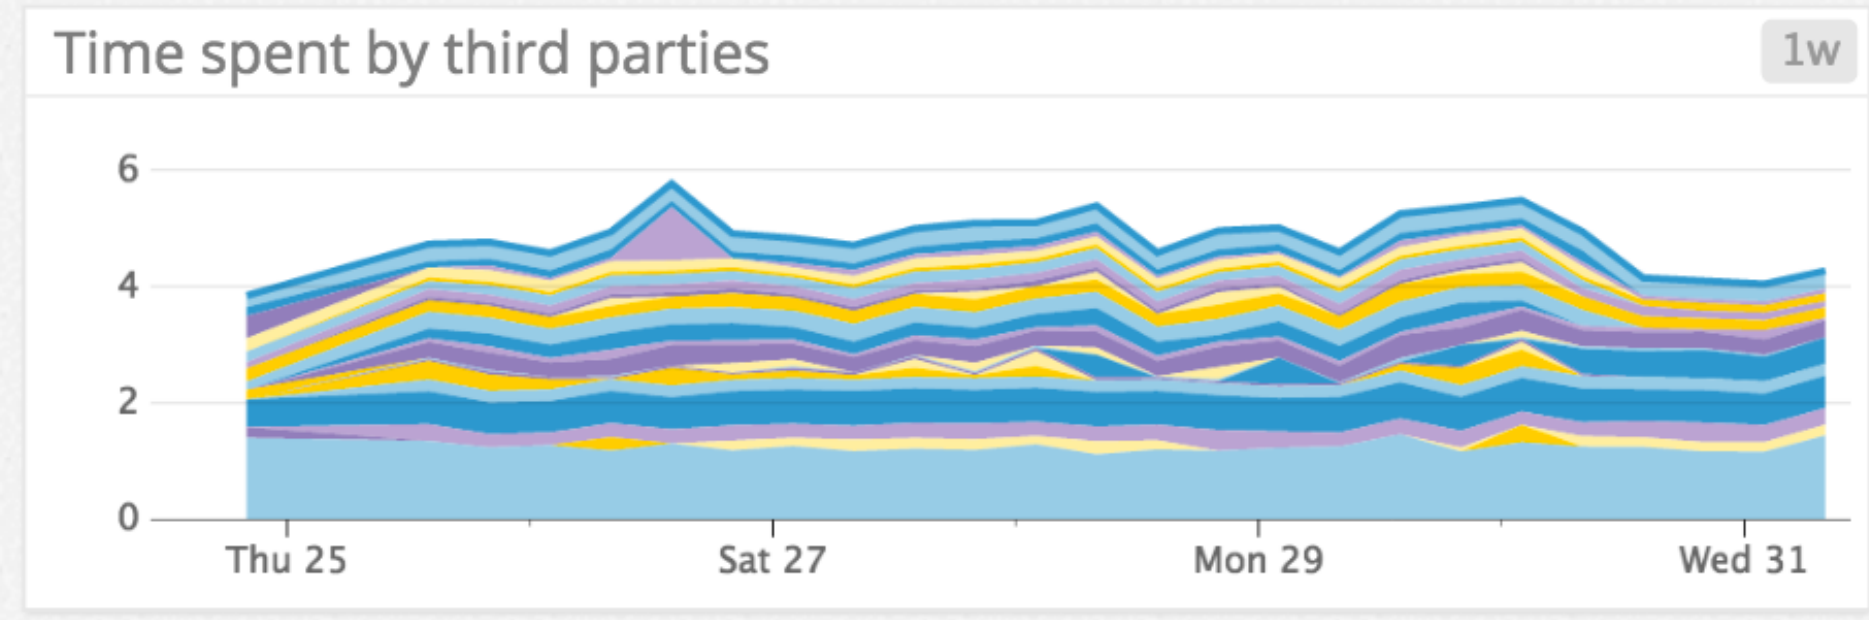


Resources



Slowest results 1w

54.05	tokyo_jp
50.84	frankfurt_de
46.53	sydney_au
43.77	fremont_ca_usa
38.95	newark_nj_usa
38.73	singapore_sg
32.3	dallas_tx_usa
28.87	são_paulo_br
18.47	san_francisco_ca_usa
16.35	dublin_ie



Network



Slowest resources 1w

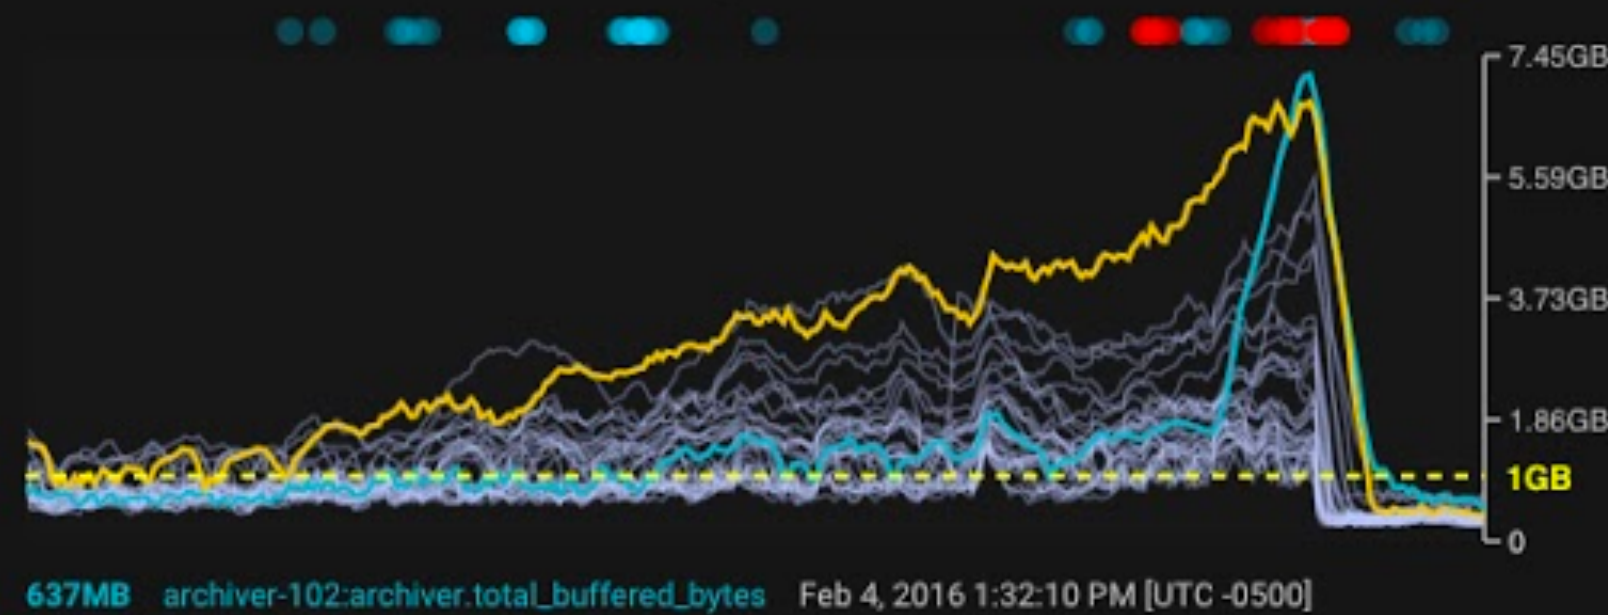
569	kissmetrics.com
480	datadoghq.com
441	optimizely.com

Overall Data Processing Health

TIME 1h 6h 1d 1w 1m 6w custom

Mon 25 Wed 27 Fri 29 Jan 31 February Wed 03 Fri 05

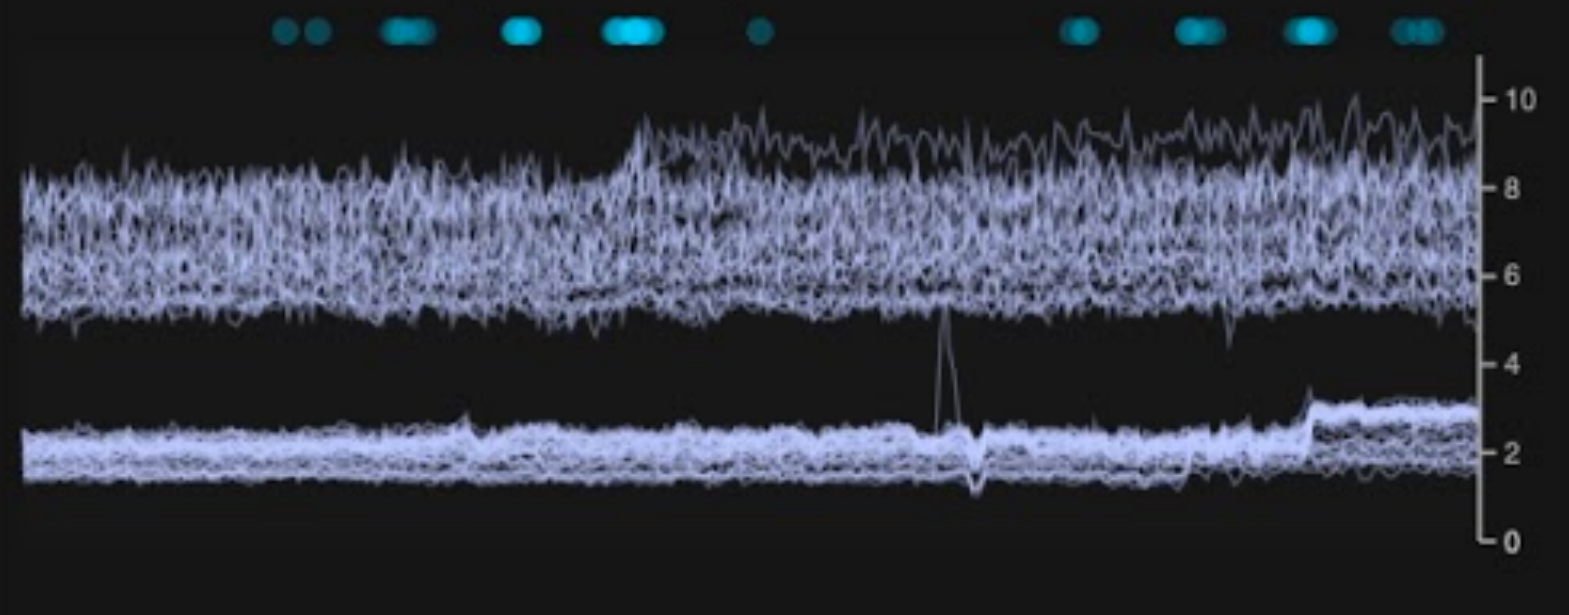
Archiver Buffer Size (Should be < 1GB)



637MB archiver-102:archiver.total_buffered_bytes Feb 4, 2016 1:32:10 PM [UTC -0500]

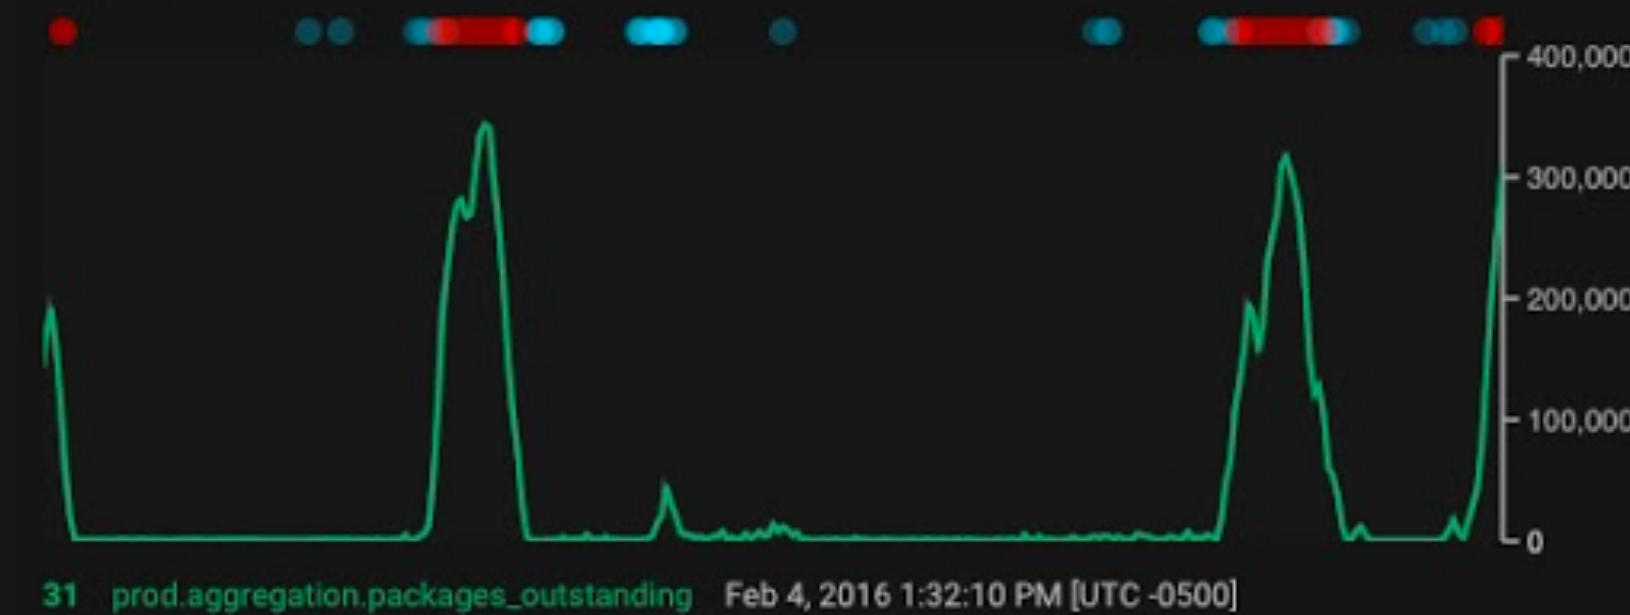
VALUE	NAME	AVERAGE	PEAK	VOLATILITY
637MB	archiver-102:archiver.total_buffered_bytes			
544MB	archiver-155:archiver.total_buffered_bytes			
680MB	archiver-146:archiver.total_buffered_bytes			
758MB	archiver-158:archiver.total_buffered_bytes			

GW Cell 5min Load Avg (Should be < 2)



VALUE	NAME	AVERAGE	PEAK	VOLATILITY
2.4	alerting-indexer-41			
1.8	alerting-indexer-43			
1.8	alerting-indexer-44			
1.9	alerting-indexer-45			

Aggregation Jobs Pending



31 prod.aggregation.packages_outstanding Feb 4, 2016 1:32:10 PM [UTC -0500]

VALUE	NAME	AVERAGE	PEAK	VOLATILITY
31	prod.aggregation.packages_outstanding			

gateway Load Average (past 5m) [Average]



VALUE	NAME	AVERAGE	PEAK	VOLATILITY
5.89	gateway-11			
6.67	gateway-12			
5.15	gateway-14			

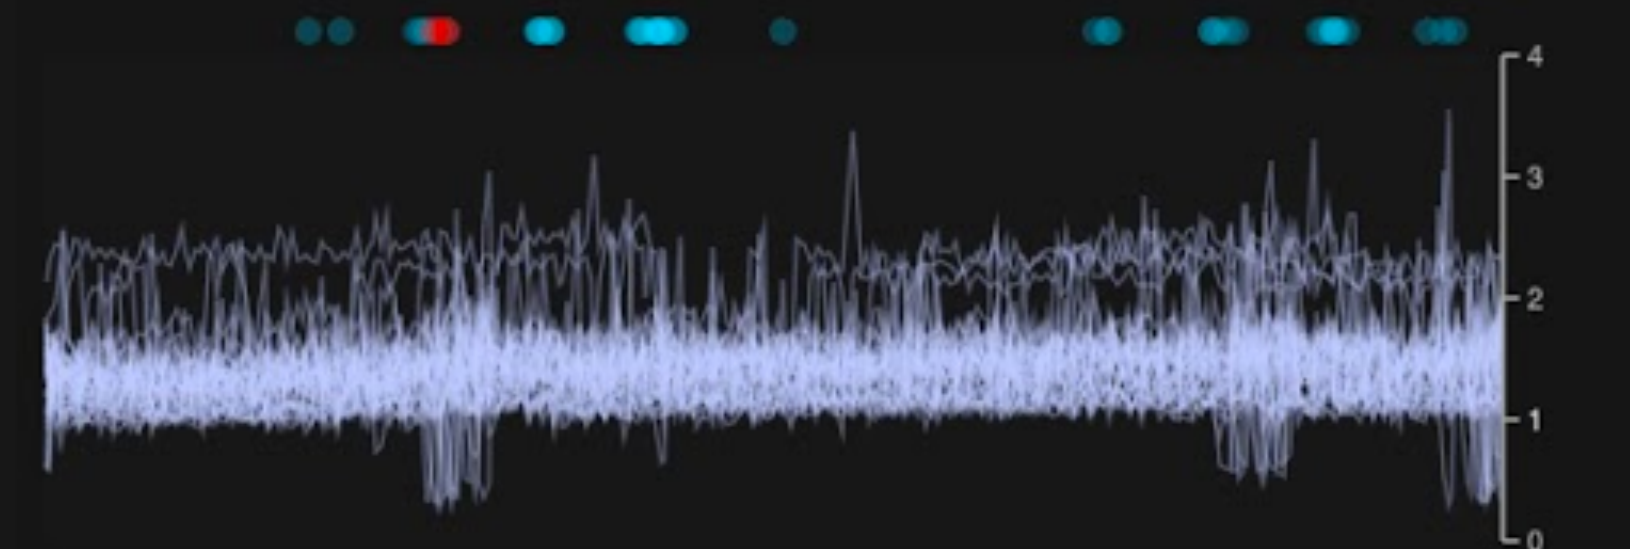
Load Balancer



2,658/s nginx requests Feb 4, 2016 1:32:10 PM [UTC -0500]

VALUE	NAME	AVERAGE	PEAK	VOLATILITY
2,658/s	nginx requests			

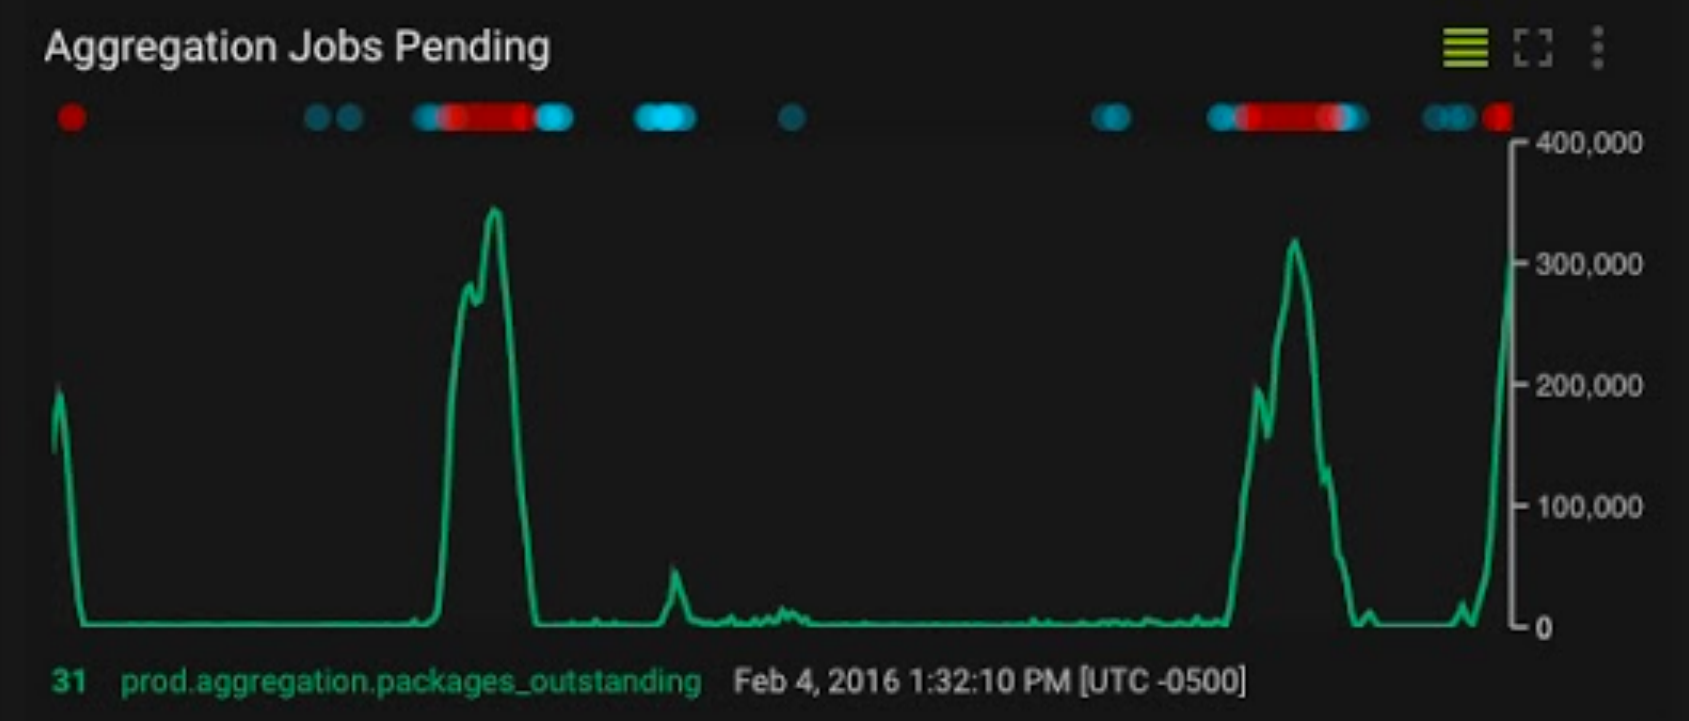
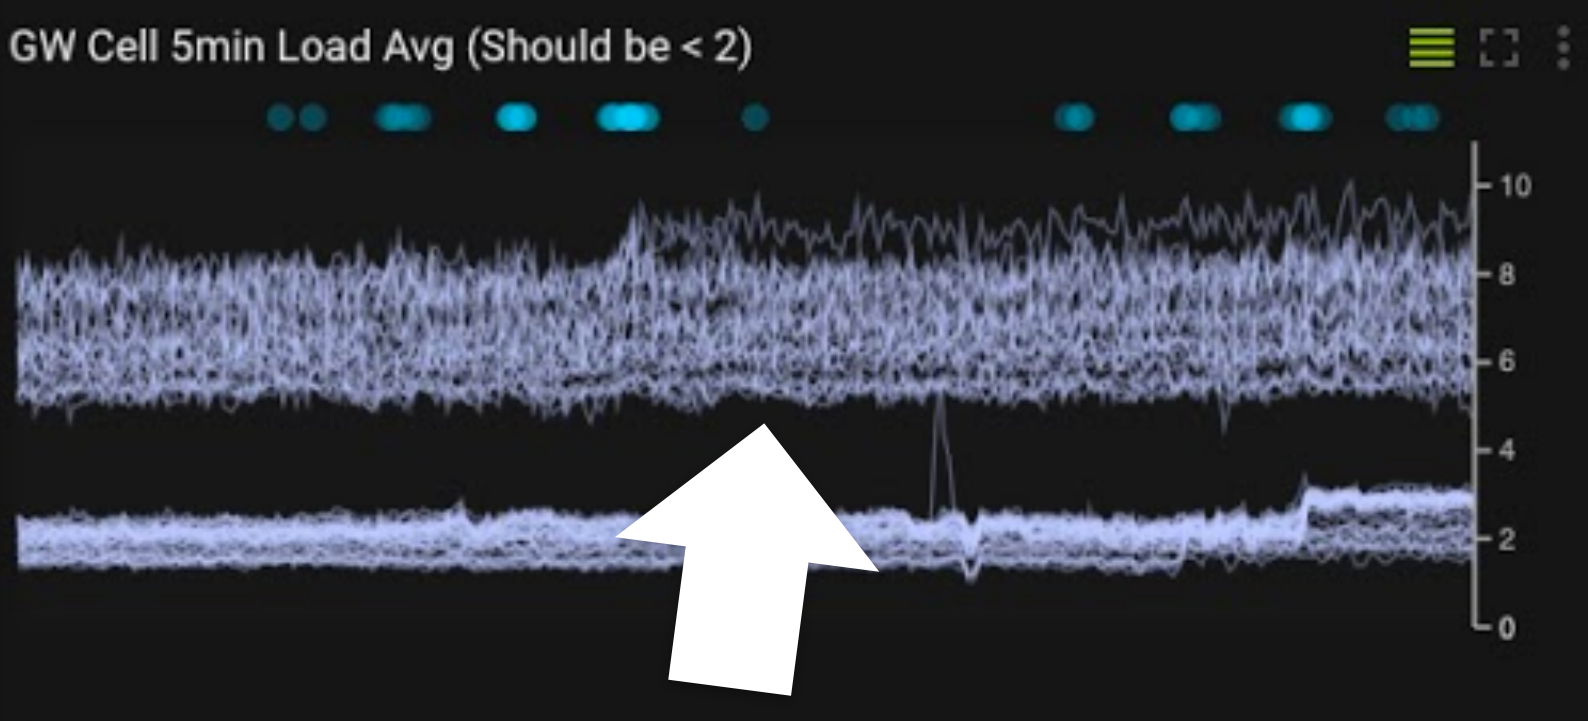
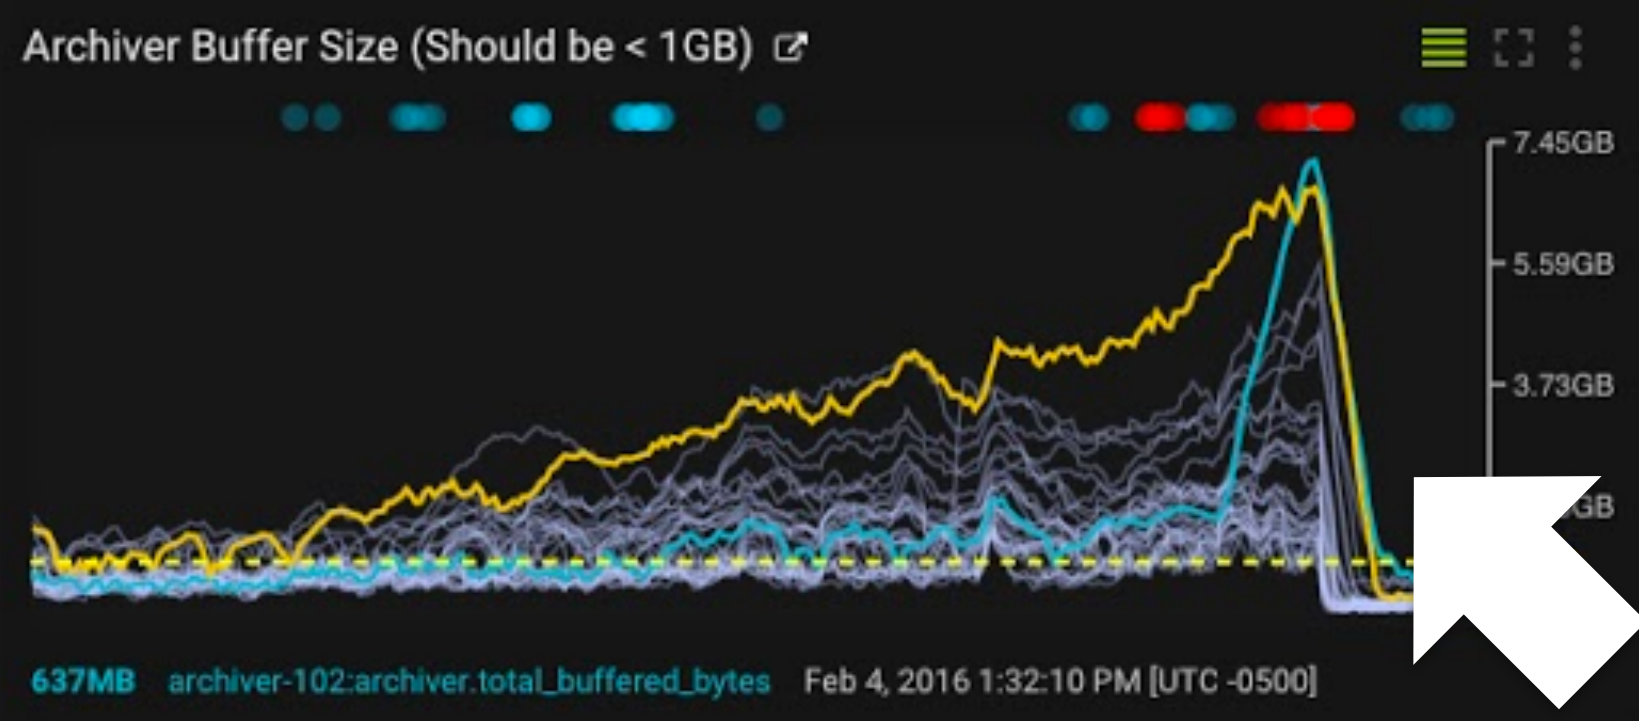
Cassandra 5min Load Avg (Should be < 15ish)



VALUE	NAME	AVERAGE	PEAK	VOLATILITY
1.41	cassandra-9			
1.27	cassandra-8			
1.41	cassandra-7			

Overall Data Processing Health

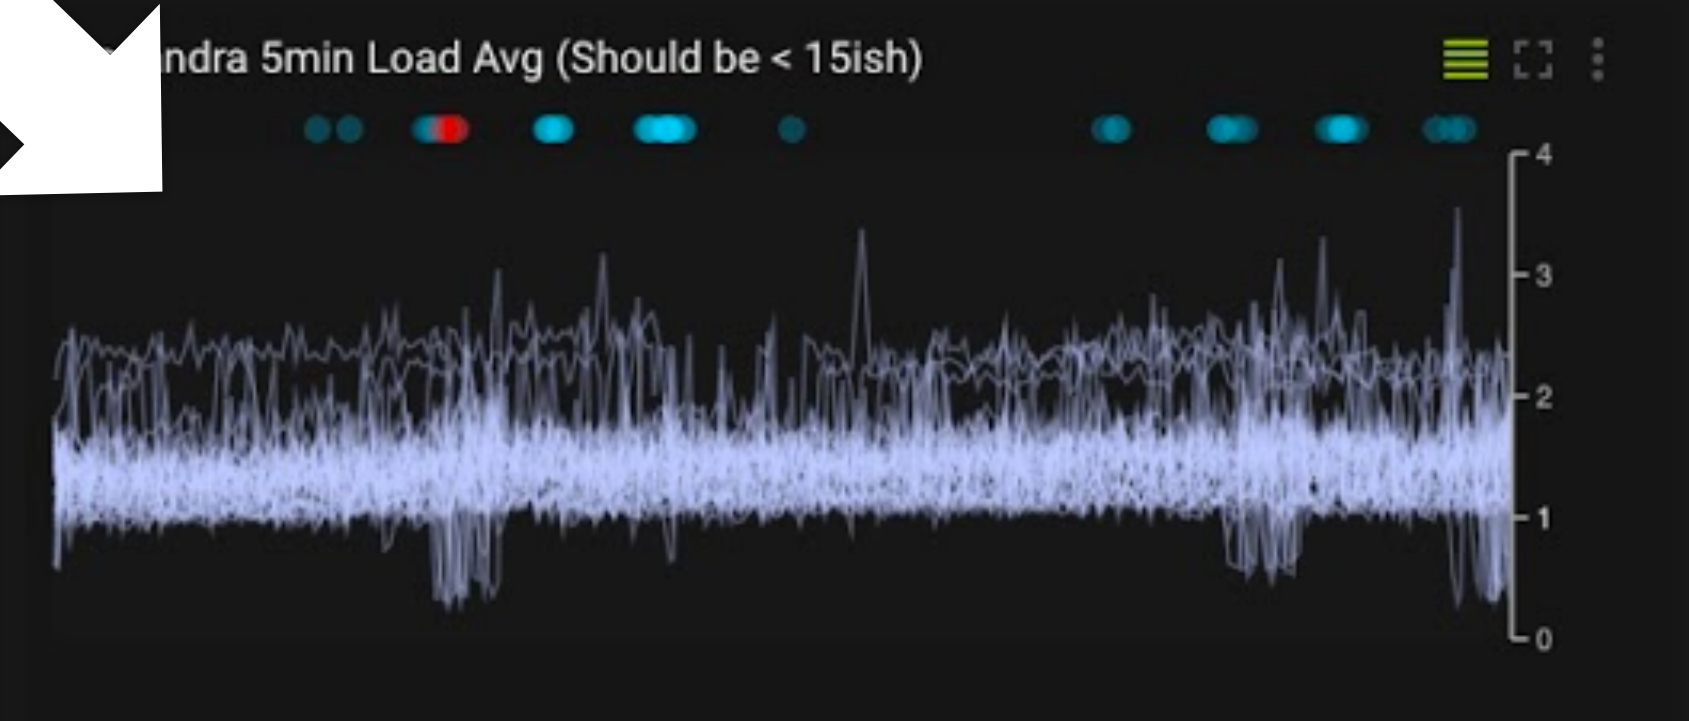
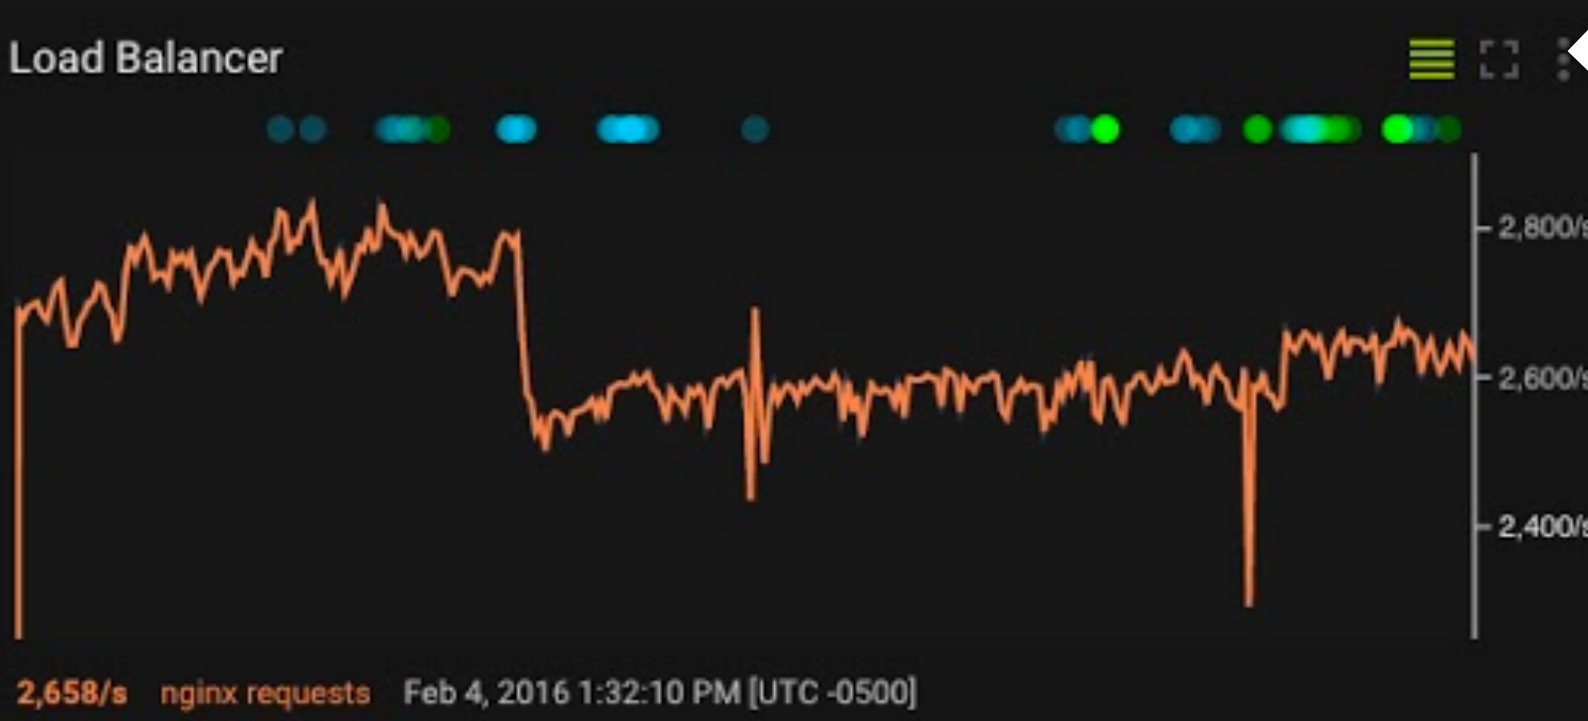
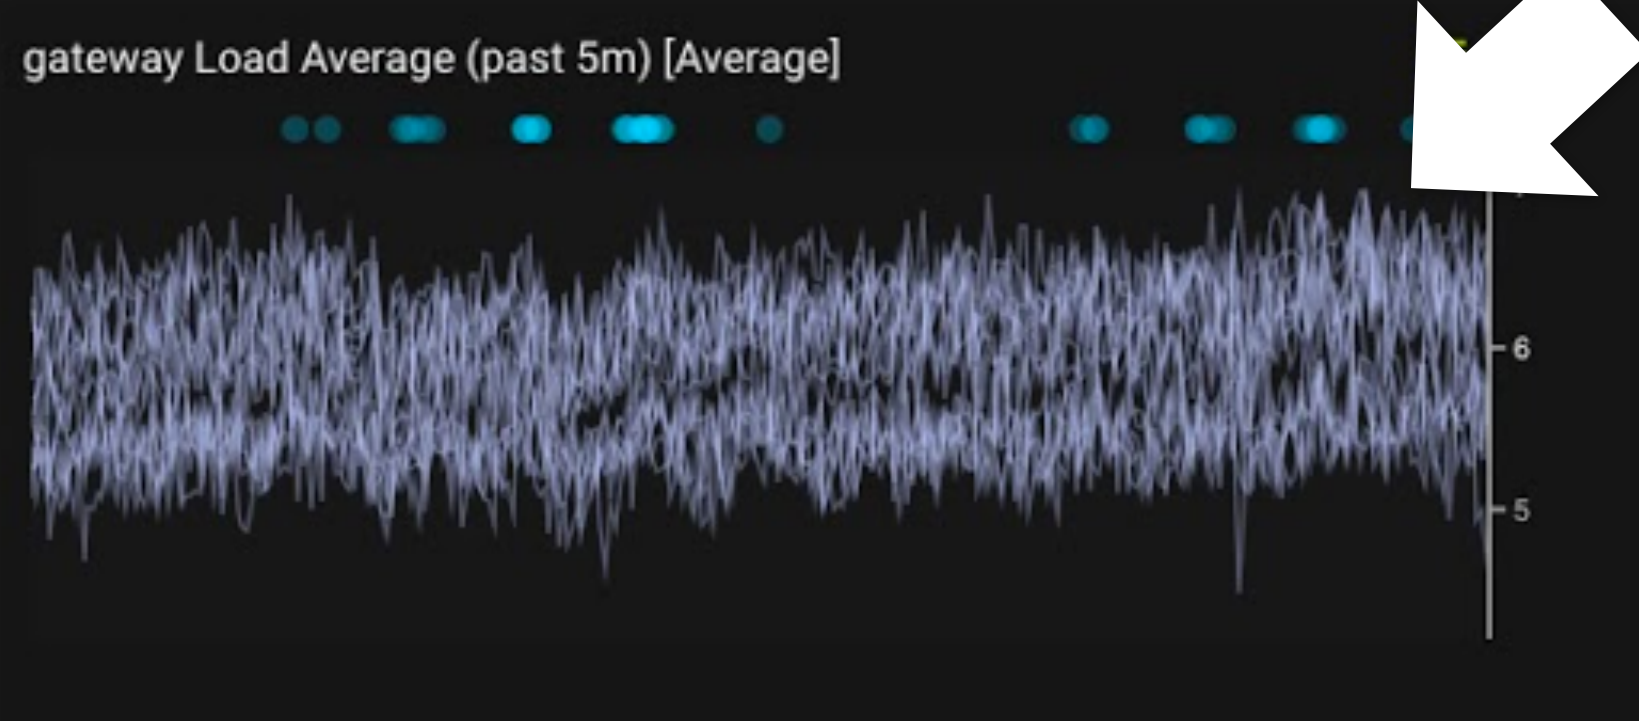
Mon 25 Wed 27 Fri 29 Jan 31 February Wed 03 Fri 05 Mon 25 Wed 27 Fri 29 Jan 31 February Wed 03 Fri 05 Mon 25 Wed 27 Fri 29 Jan 31 February Wed 03 Fri 05



VALUE	NAME	AVERAGE
637MB	archiver-102:archiver.total_buffered_bytes	
544MB	archiver-155:archiver.total_buffered_bytes	
680MB	archiver-146:archiver.total_buffered_bytes	
758MB	archiver-158:archiver.total_buffered_bytes	

different dimensions of X splattered across the screen

	AVERAGE	PEAK	VOLATILITY
prod.aggregation.packages_outstanding			



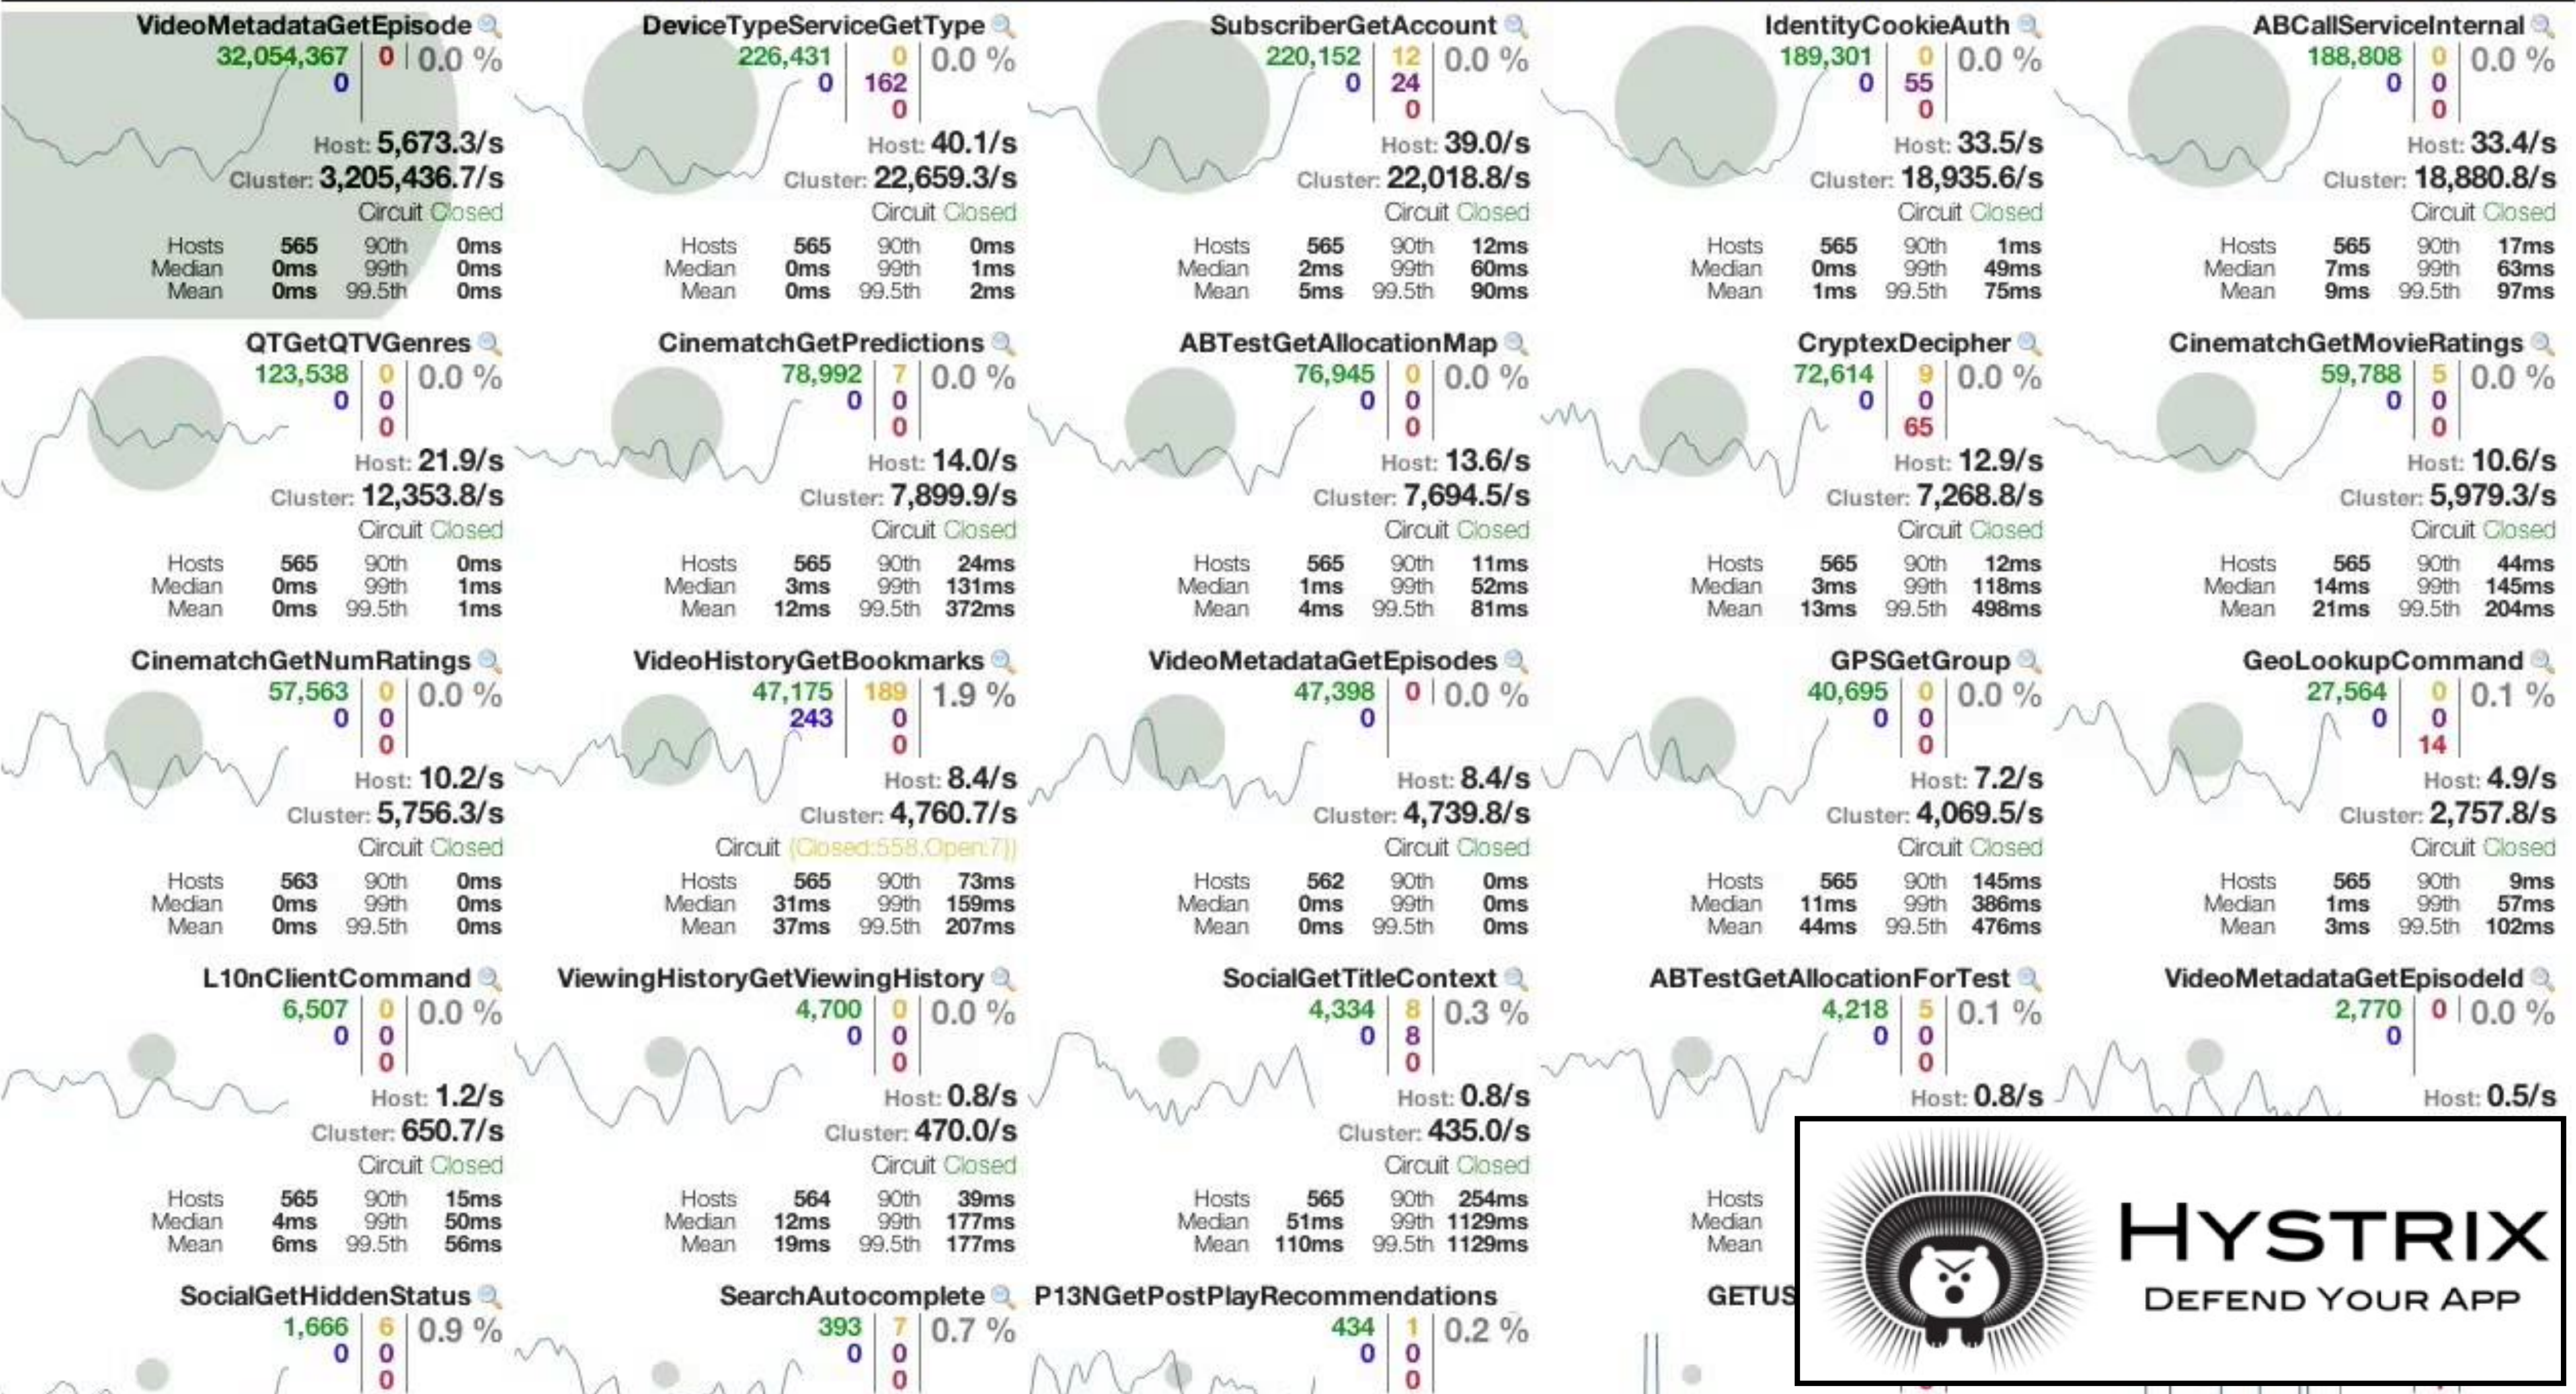
VALUE	NAME	AVERAGE	PEAK	VOLATILITY
5.89	gateway-11			
6.67	gateway-12			
5.15	gateway-14			

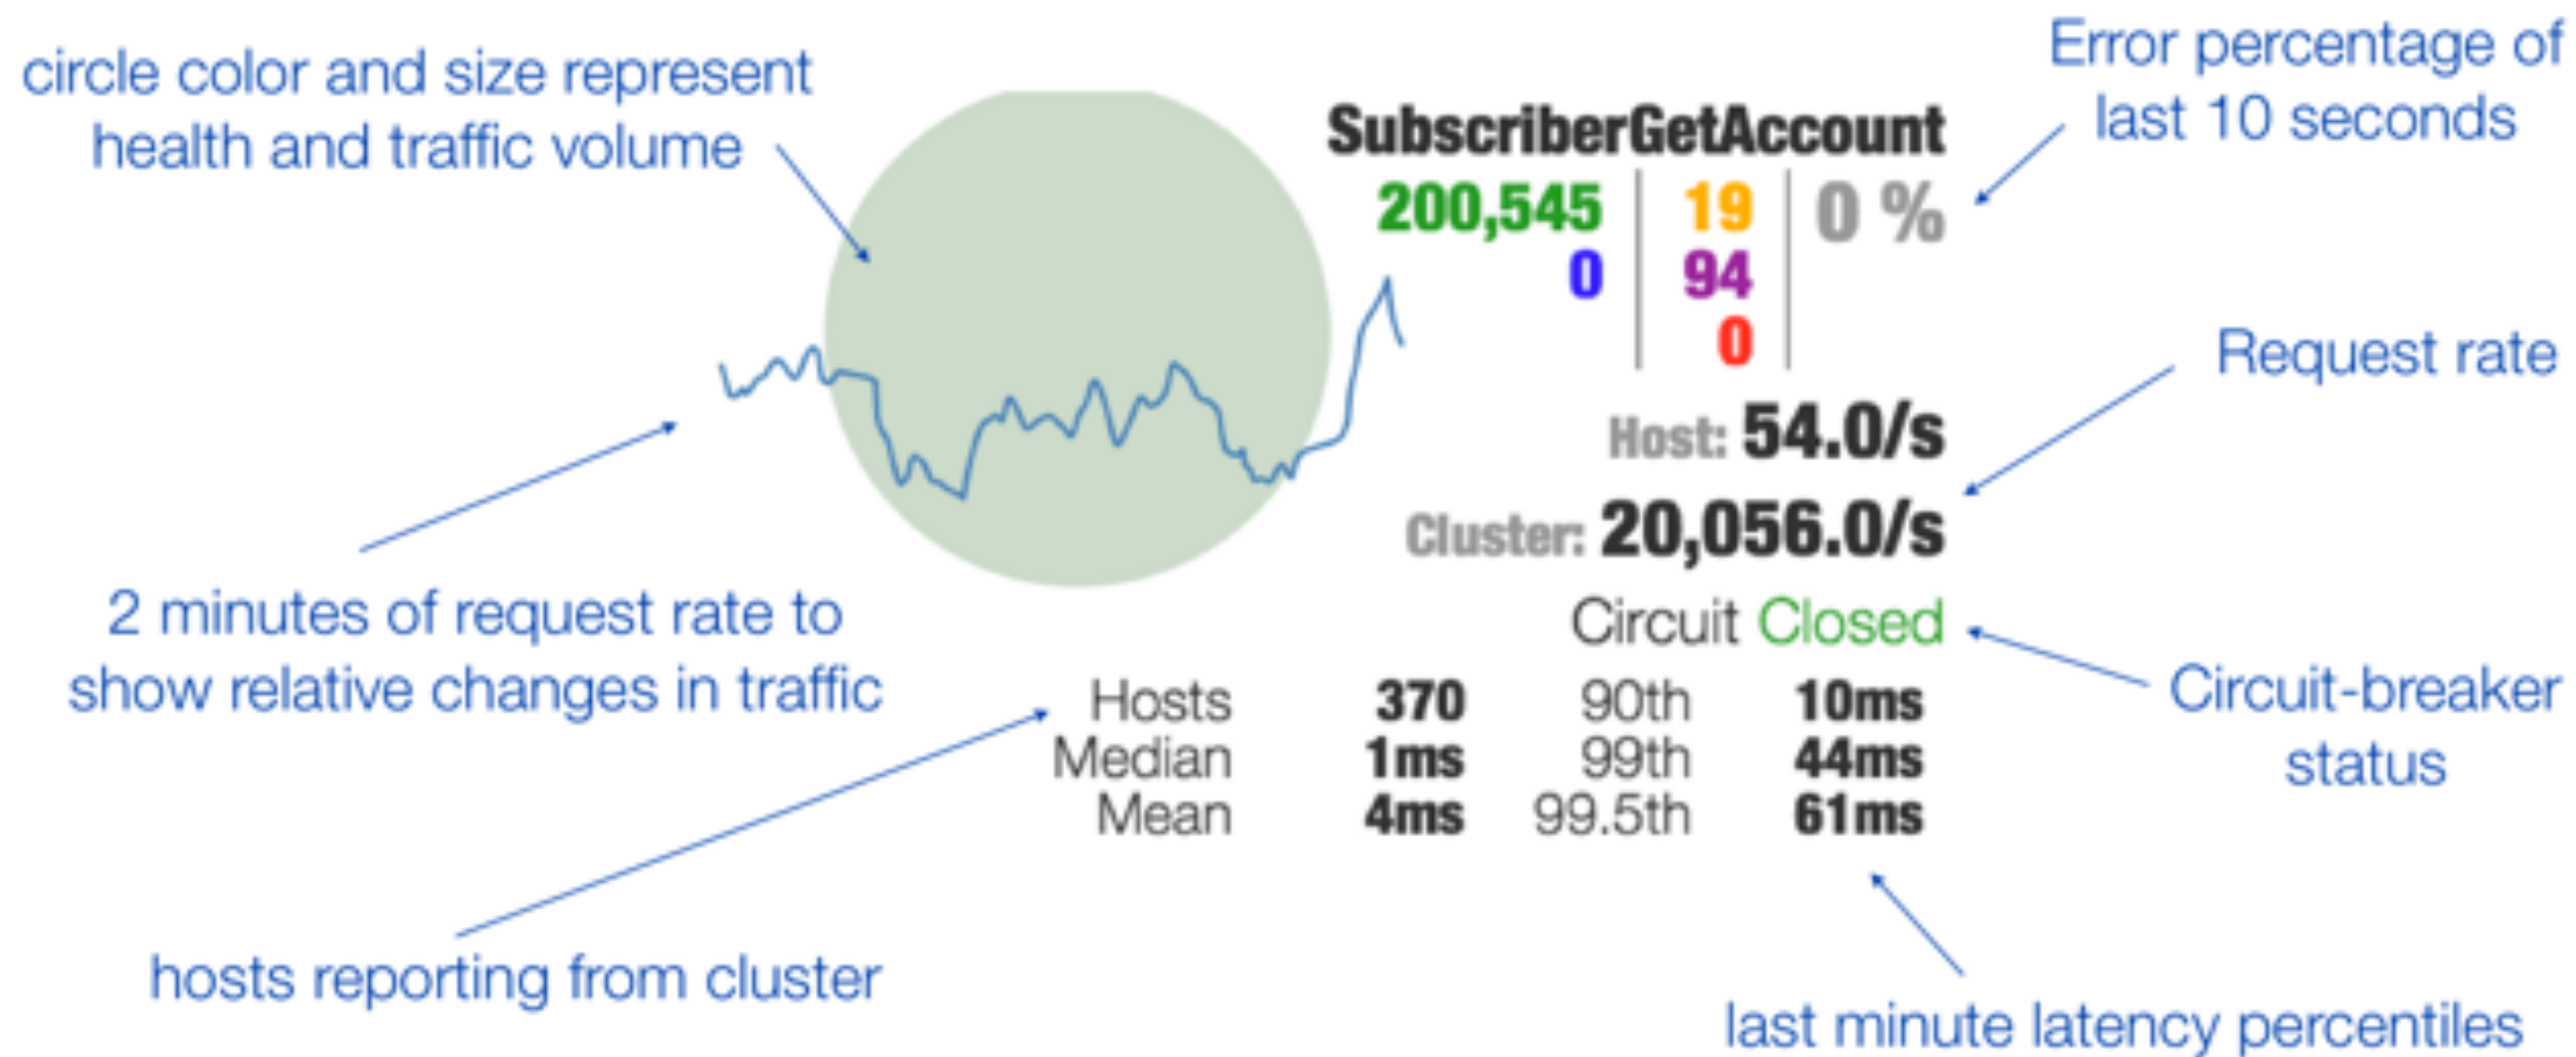
VALUE	NAME	AVERAGE	PEAK	VOLATILITY
2,658/s	nginx requests			

VALUE	NAME	AVERAGE	PEAK	VOLATILITY
1.41	cassandra-9			
1.27	cassandra-8			
1.41	cassandra-7			

Circuit Breakers

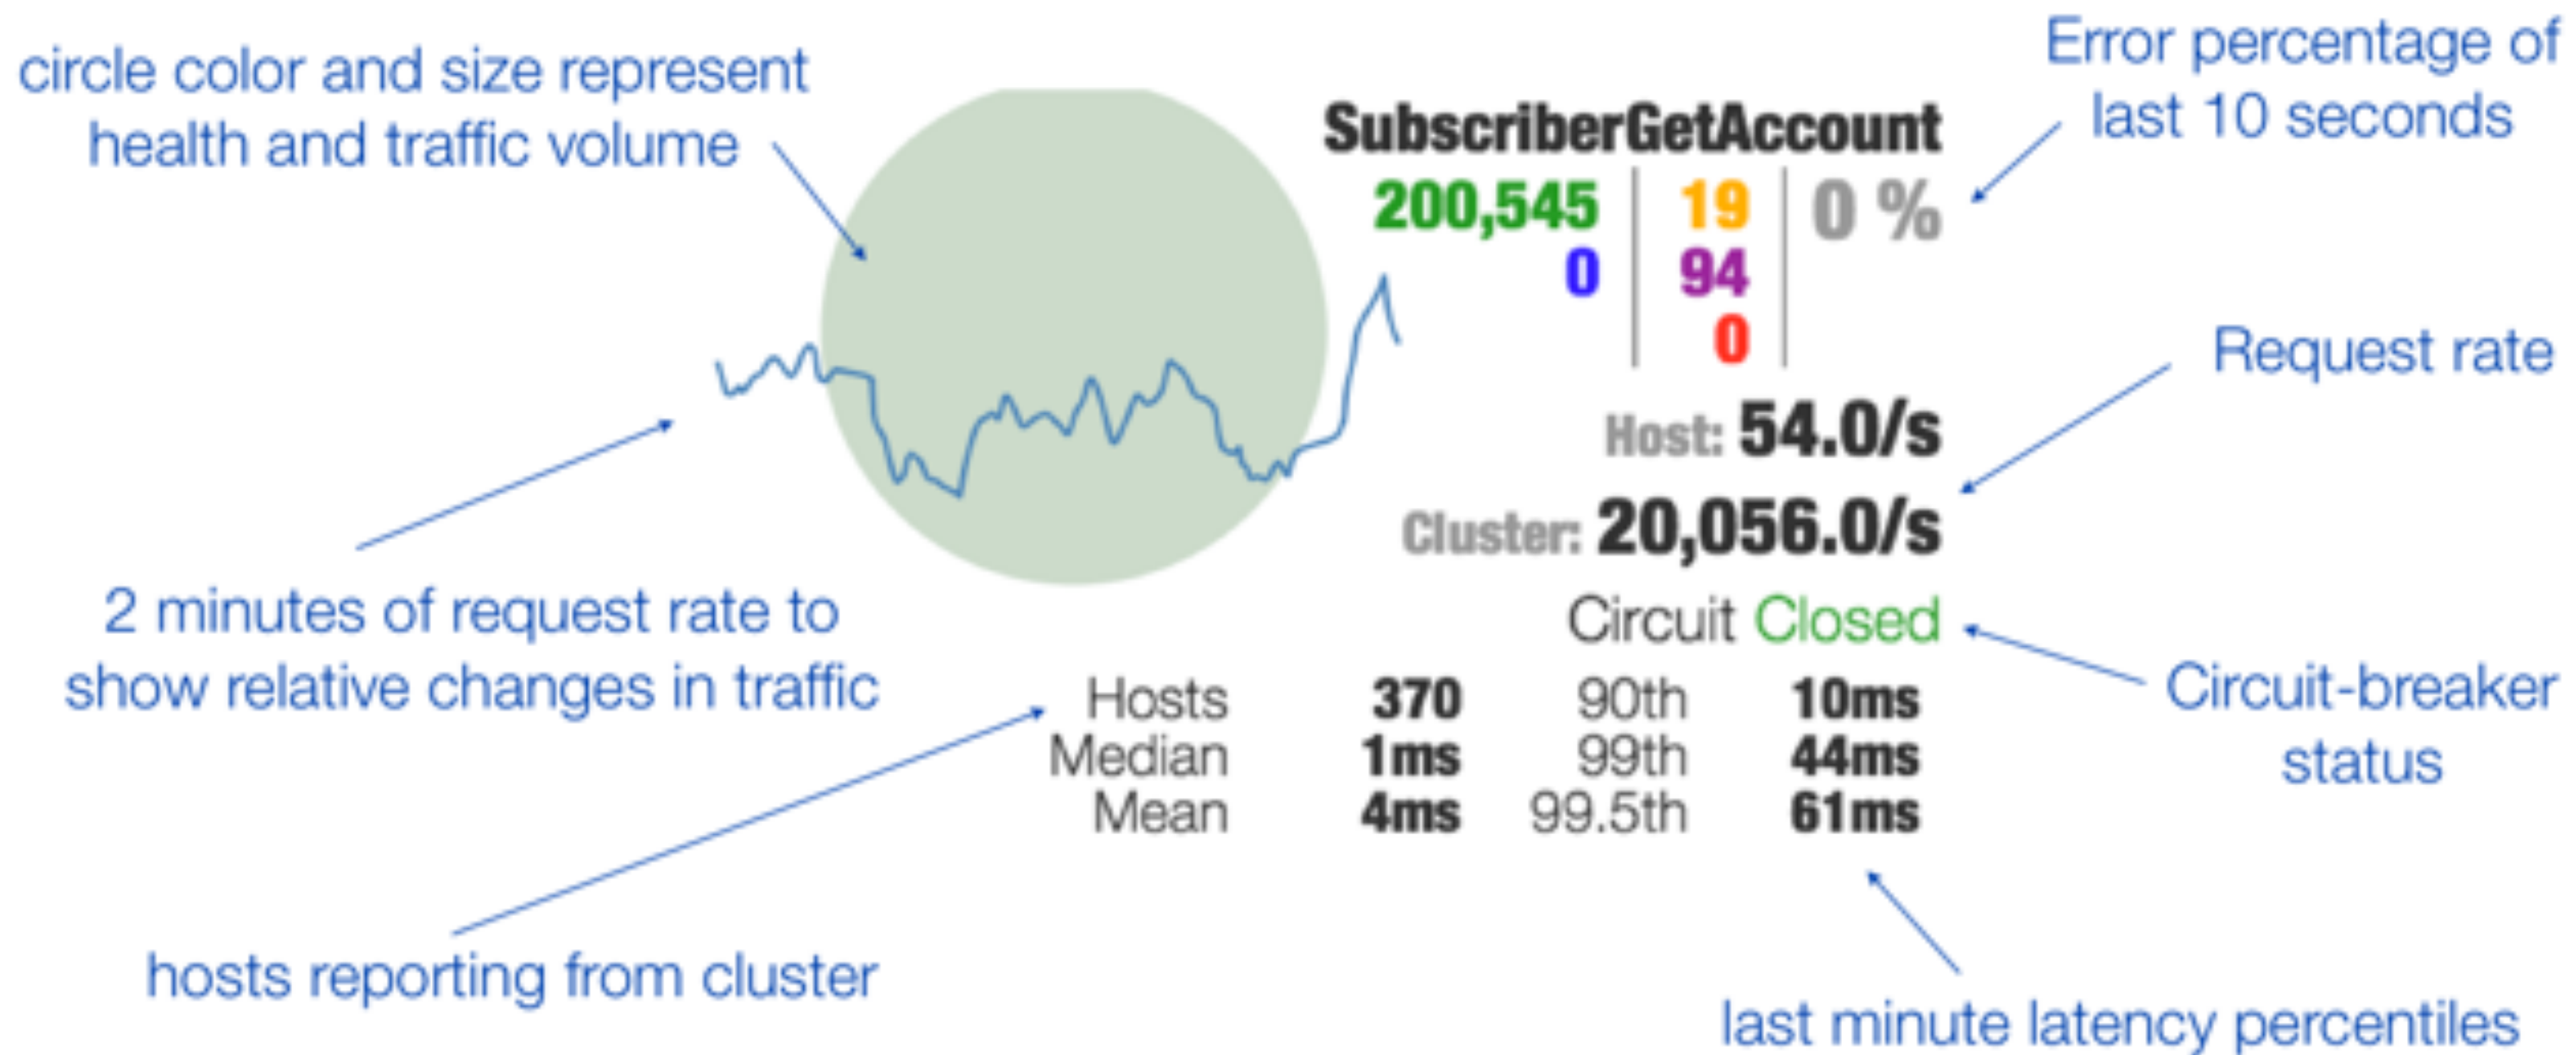
Sort: [Error then Volume](#) | [Alphabetical](#) | [Volume](#) | [Error](#) | [Mean](#) | [Median](#) | [90](#) | [99](#) | [99.5](#) | [Success](#) | [Latent](#) | [Short-Circuited](#) | [Timeout](#) | [Rejected](#) | [Failure](#) | [Error %](#)





Rolling 10 second counters
with 1 second granularity

Successes	200,545	19	Thread timeouts
Short-circuited (rejected)	0	94	Thread-pool Rejections
		0	Failures/Exceptions

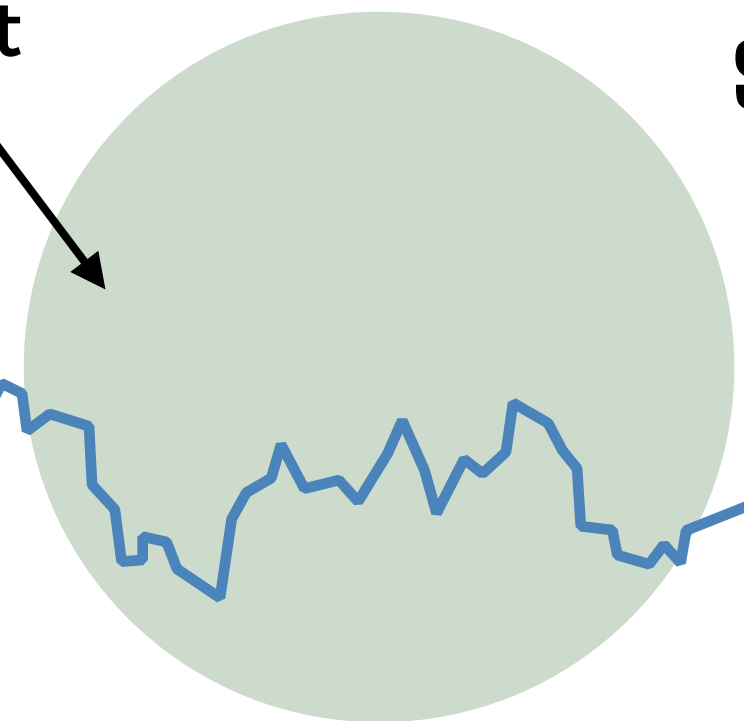


Rolling 10 second counters with 1 second granularity

Successes	200,545	19	Thread time
Short-circuited (rejected)	0	94	Thread-pool
		0	Failures/Exc

- + cold starts
- + throttled invocations
- + concurrent executions
- + estimated cost (\$)

circle colour and size represent health and traffic volume



SubscriberGetAccount

200,545 | 19 | 0 %
0 | 94 | 0 %

Est Cost: \$54.0/s

Req Rate: 20,056.0/s

Error percentage of last 10 seconds

Cold start percentage last 10 seconds

Estimated cost

Request rate

2 minutes of request rate to show relative changes in traffic

no. of concurrent executions of this function

Concurrency
Median
Mean

370
1ms
4ms

90th
99th
99.5th

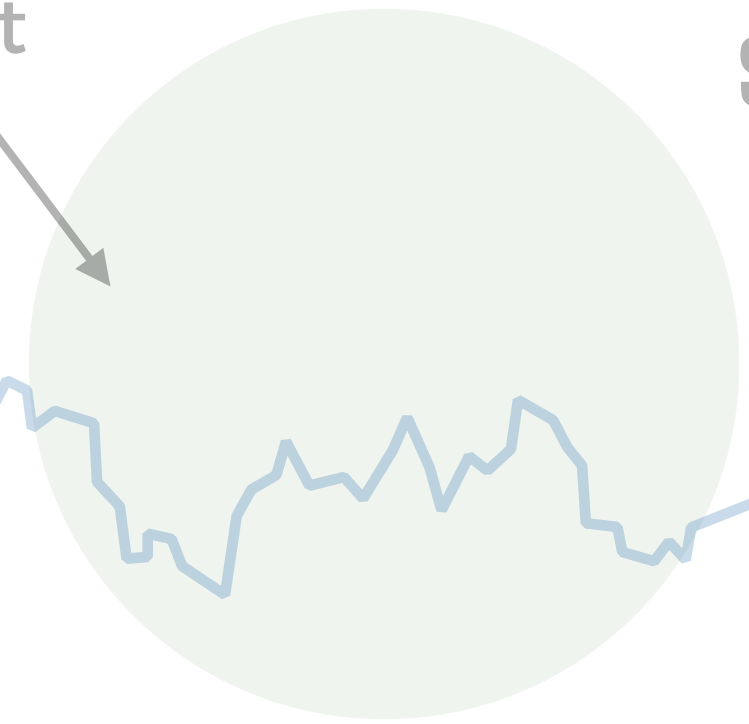
10ms
44ms
61ms

last minute latency percentiles

Rolling 10 second counters with 1 second granularity

Successes 200,545 | 19 Timeouts
Cold starts 0 | 94 Errors
0 Throttled Invocations

circle colour and size represent health and traffic volume



2 minutes of request rate to show relative changes in traffic

no. of concurrent executions of this function

Concurrency
Median
Mean

370 90th
1ms 99th
4ms 99.5th

10ms
44ms
61ms

last minute latency percentiles

SubscriberGetAccount

200,545

19

0 %

0

94

0 %

Est Cost: \$54.0/s

Req Rate: 20,056.0/s

Error percentage of last 10 seconds

Cold start percentage last 10 seconds

Estimated cost

Request rate

Rolling 10 second counters with 1 second granularity

Successes 200,545

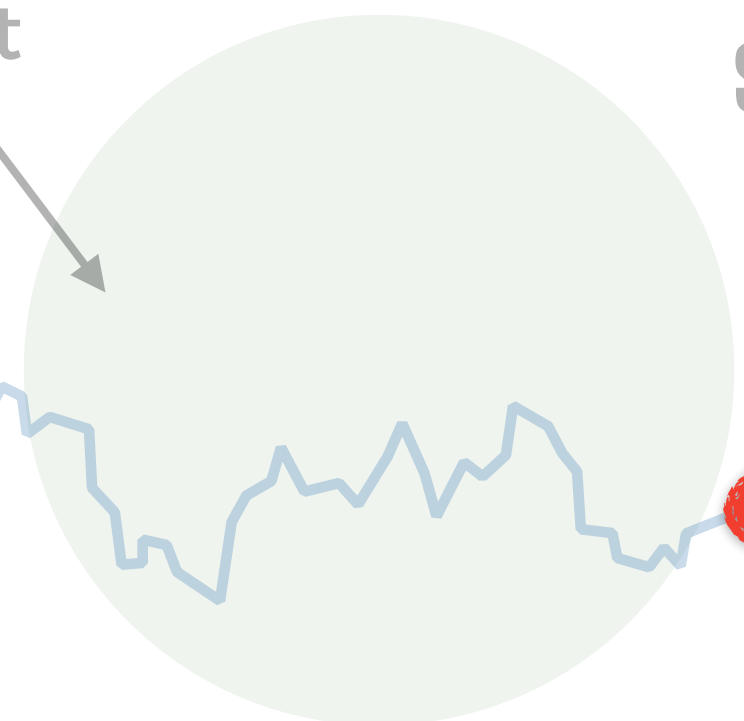
19 Timeouts

Cold starts 0

94 Errors

0 Throttled Invocations

circle colour and size represent health and traffic volume



SubscriberGetAccount

200,545
0 | 19 | 0 %
0 | 94 | 0 %

Error percentage of last 10 seconds

Cold start percentage last 10 seconds

Est Cost: **\$54.0/s**

Estimated cost

Req Rate: 20,056.0/s

Request rate

2 minutes of request rate to show relative changes in traffic

Concurrency
Median
Mean

370
1ms
4ms | 90th
99th
99.5th | 10ms
44ms
61ms

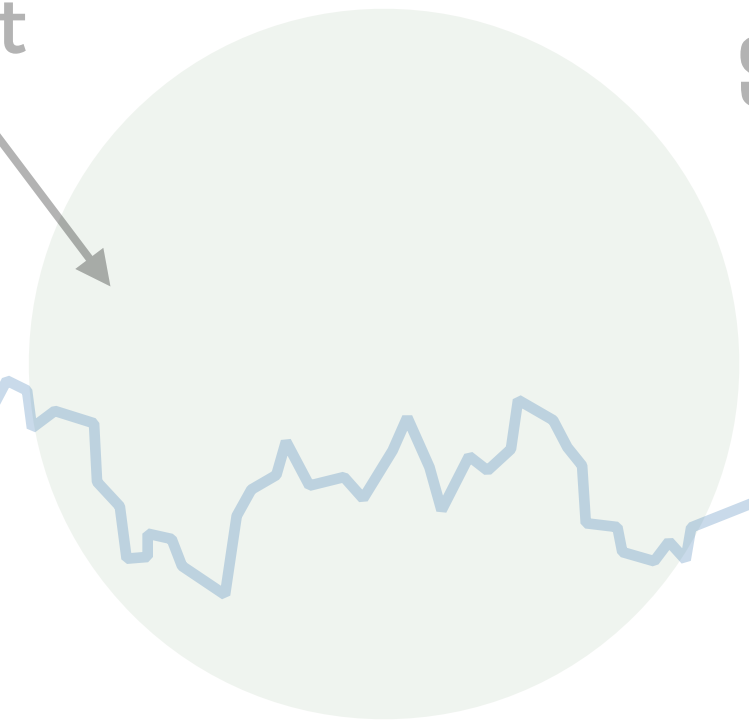
no. of concurrent executions of this function

last minute latency percentiles

Rolling 10 second counters with 1 second granularity

Successes 200,545 | 19 Timeouts
Cold starts 0 | 94 Errors
0 Throttled Invocations

circle colour and size represent health and traffic volume



SubscriberGetAccount

200,545 | 19 | 0 %
0 | 94 | 0 %

Est Cost: \$54.0/s

Req Rate: 20,056.0/s

Error percentage of last 10 seconds

Cold start percentage last 10 seconds

Estimated cost

Request rate

2 minutes of request rate to show relative changes in traffic

Concurrency **370**

Median
Mean

370
1ms
4ms

90th
99th
99.5th

10ms
44ms
61ms

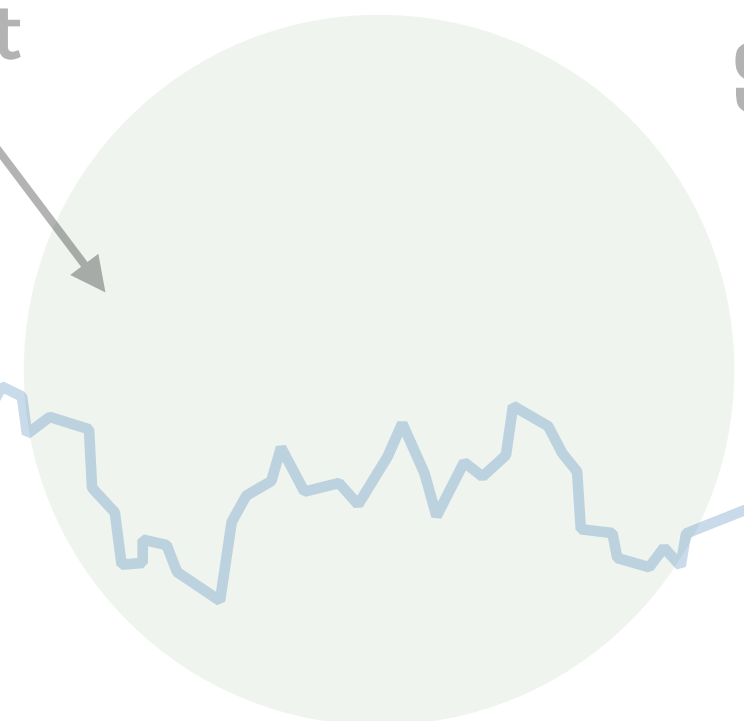
no. of concurrent executions of this function

last minute latency percentiles

Rolling 10 second counters with 1 second granularity

Successes 200,545 | 19 Timeouts
Cold starts 0 | 94 Errors
0 Throttled Invocations

circle colour and size represent health and traffic volume



SubscriberGetAccount

200,545 | 19 | 0 %
0 | 94 | 0 %

Est Cost: \$54.0/s

Req Rate: 20,056.0/s

Error percentage of last 10 seconds

Cold start percentage last 10 seconds

Estimated cost

Request rate

2 minutes of request rate to show relative changes in traffic

Concurrency
Median
Mean

370
1ms
4ms

90th
99th
99.5th

10ms
44ms
61ms

no. of concurrent executions of this function

last minute latency percentiles

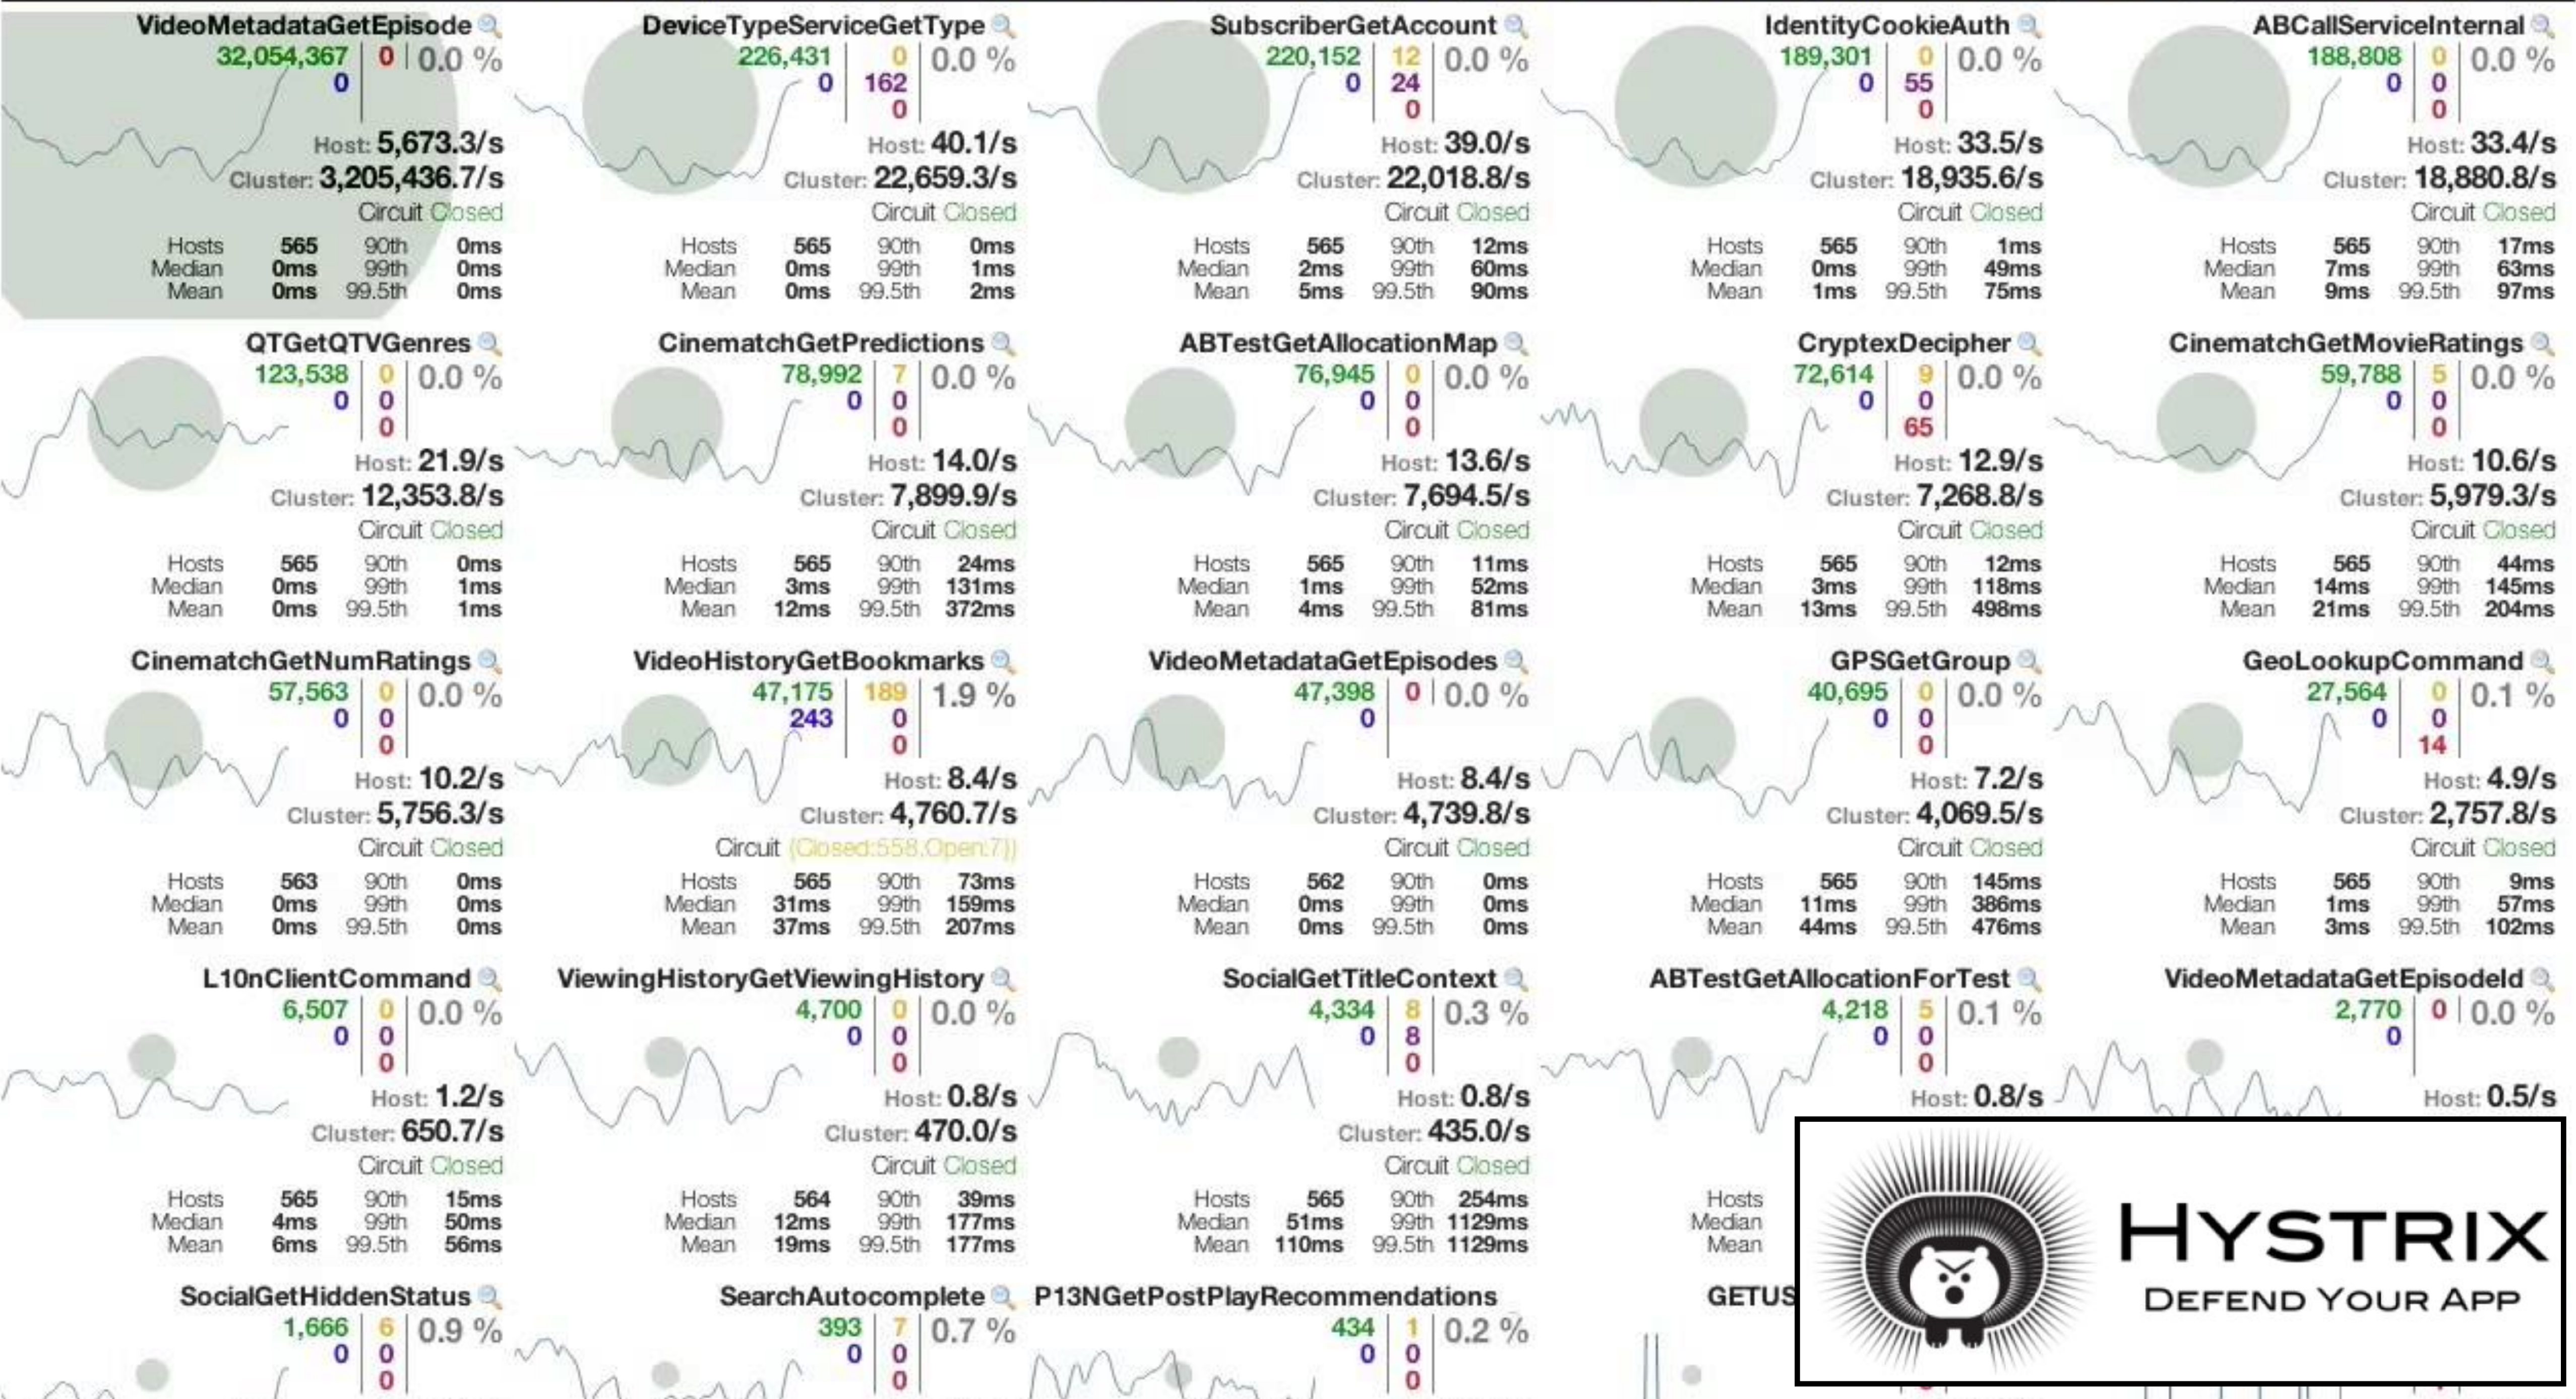
Rolling 10 second counters with 1 second granularity

Successes 200,545 | 19 Timeouts
Cold starts 0 | 94 Errors

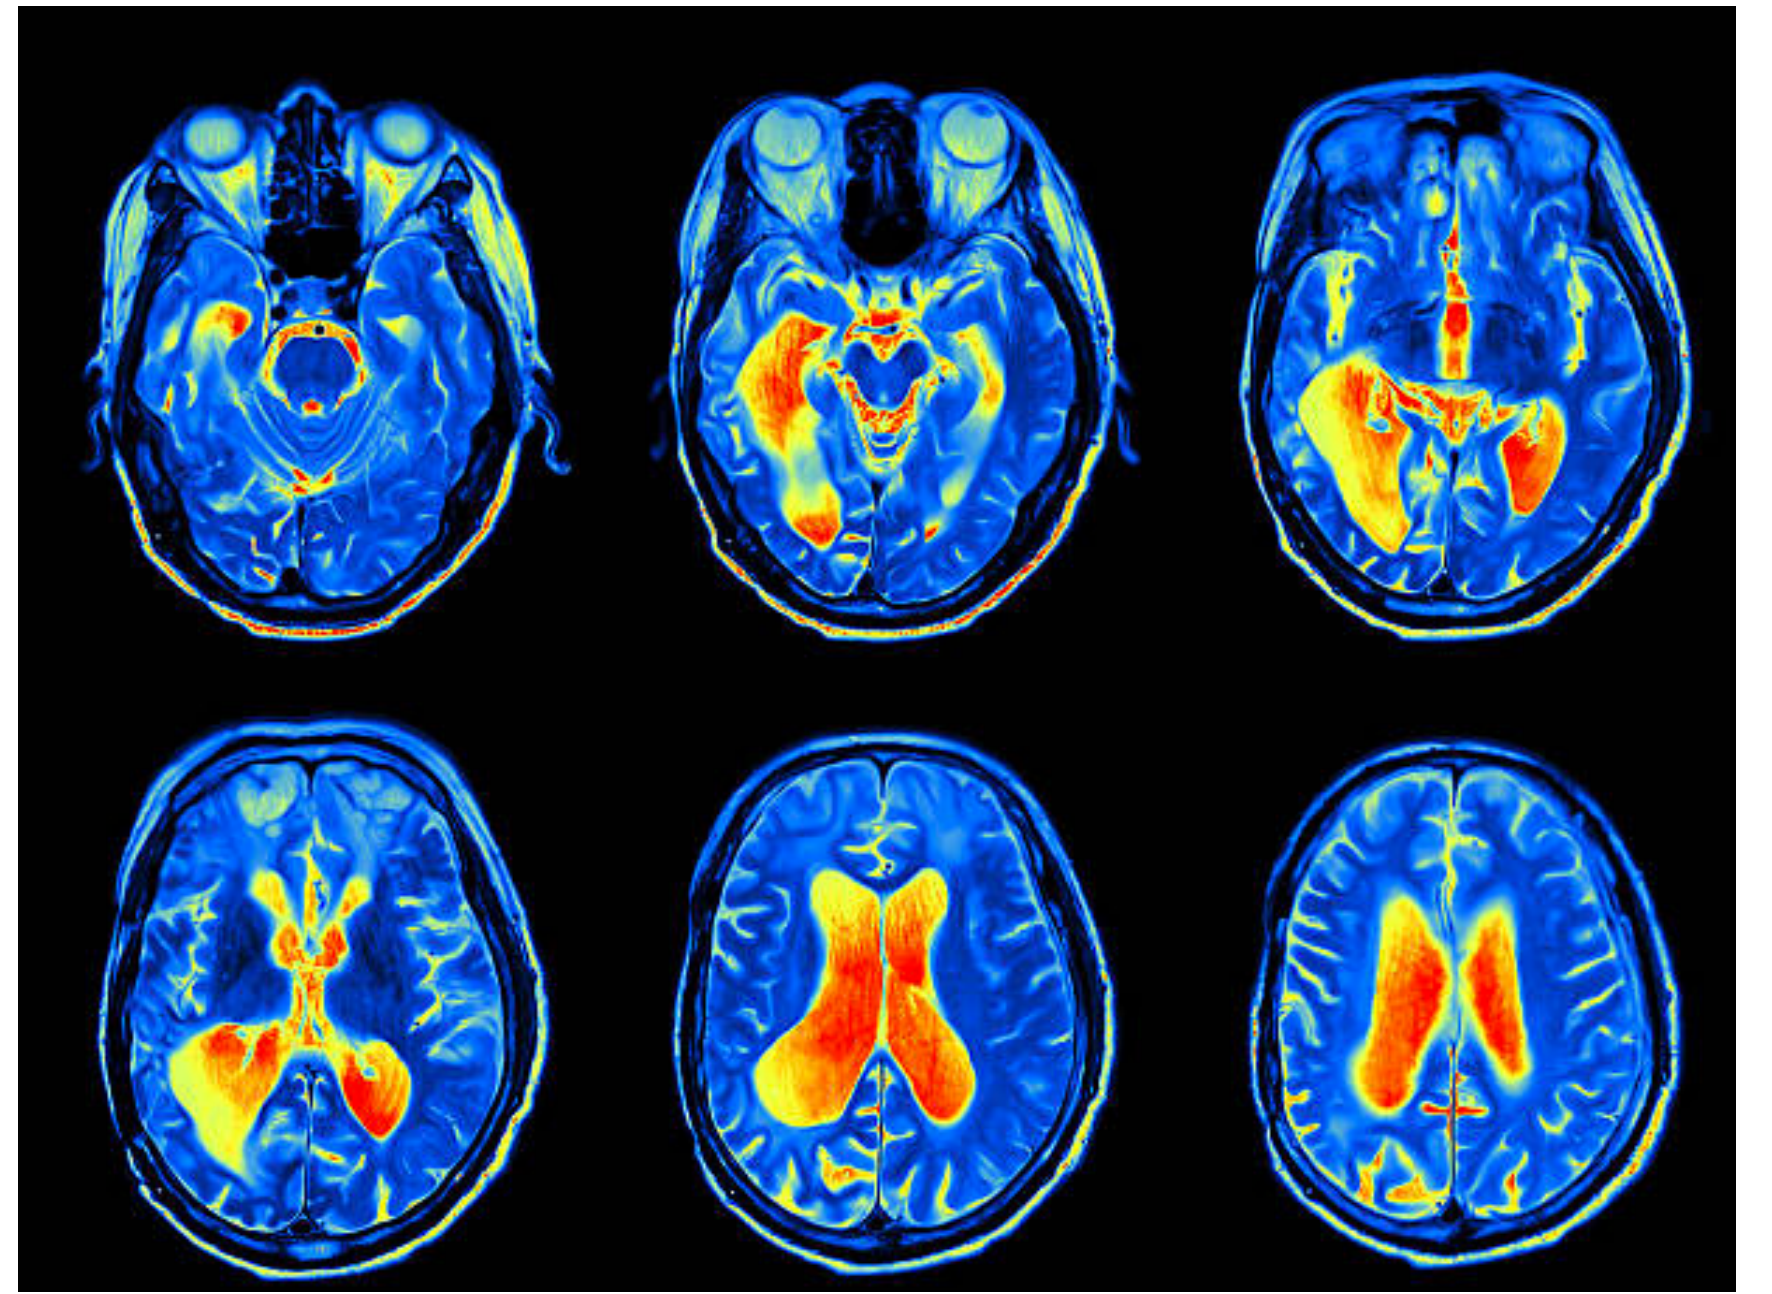
0 Throttled Invocations

Circuit Breakers

Sort: [Error then Volume](#) | [Alphabetical](#) | [Volume](#) | [Error](#) | [Mean](#) | [Median](#) | [90](#) | [99](#) | [99.5](#) | [Success](#) | [Latent](#) | [Short-Circuited](#) | [Timeout](#) | [Rejected](#) | [Failure](#) | [Error %](#)



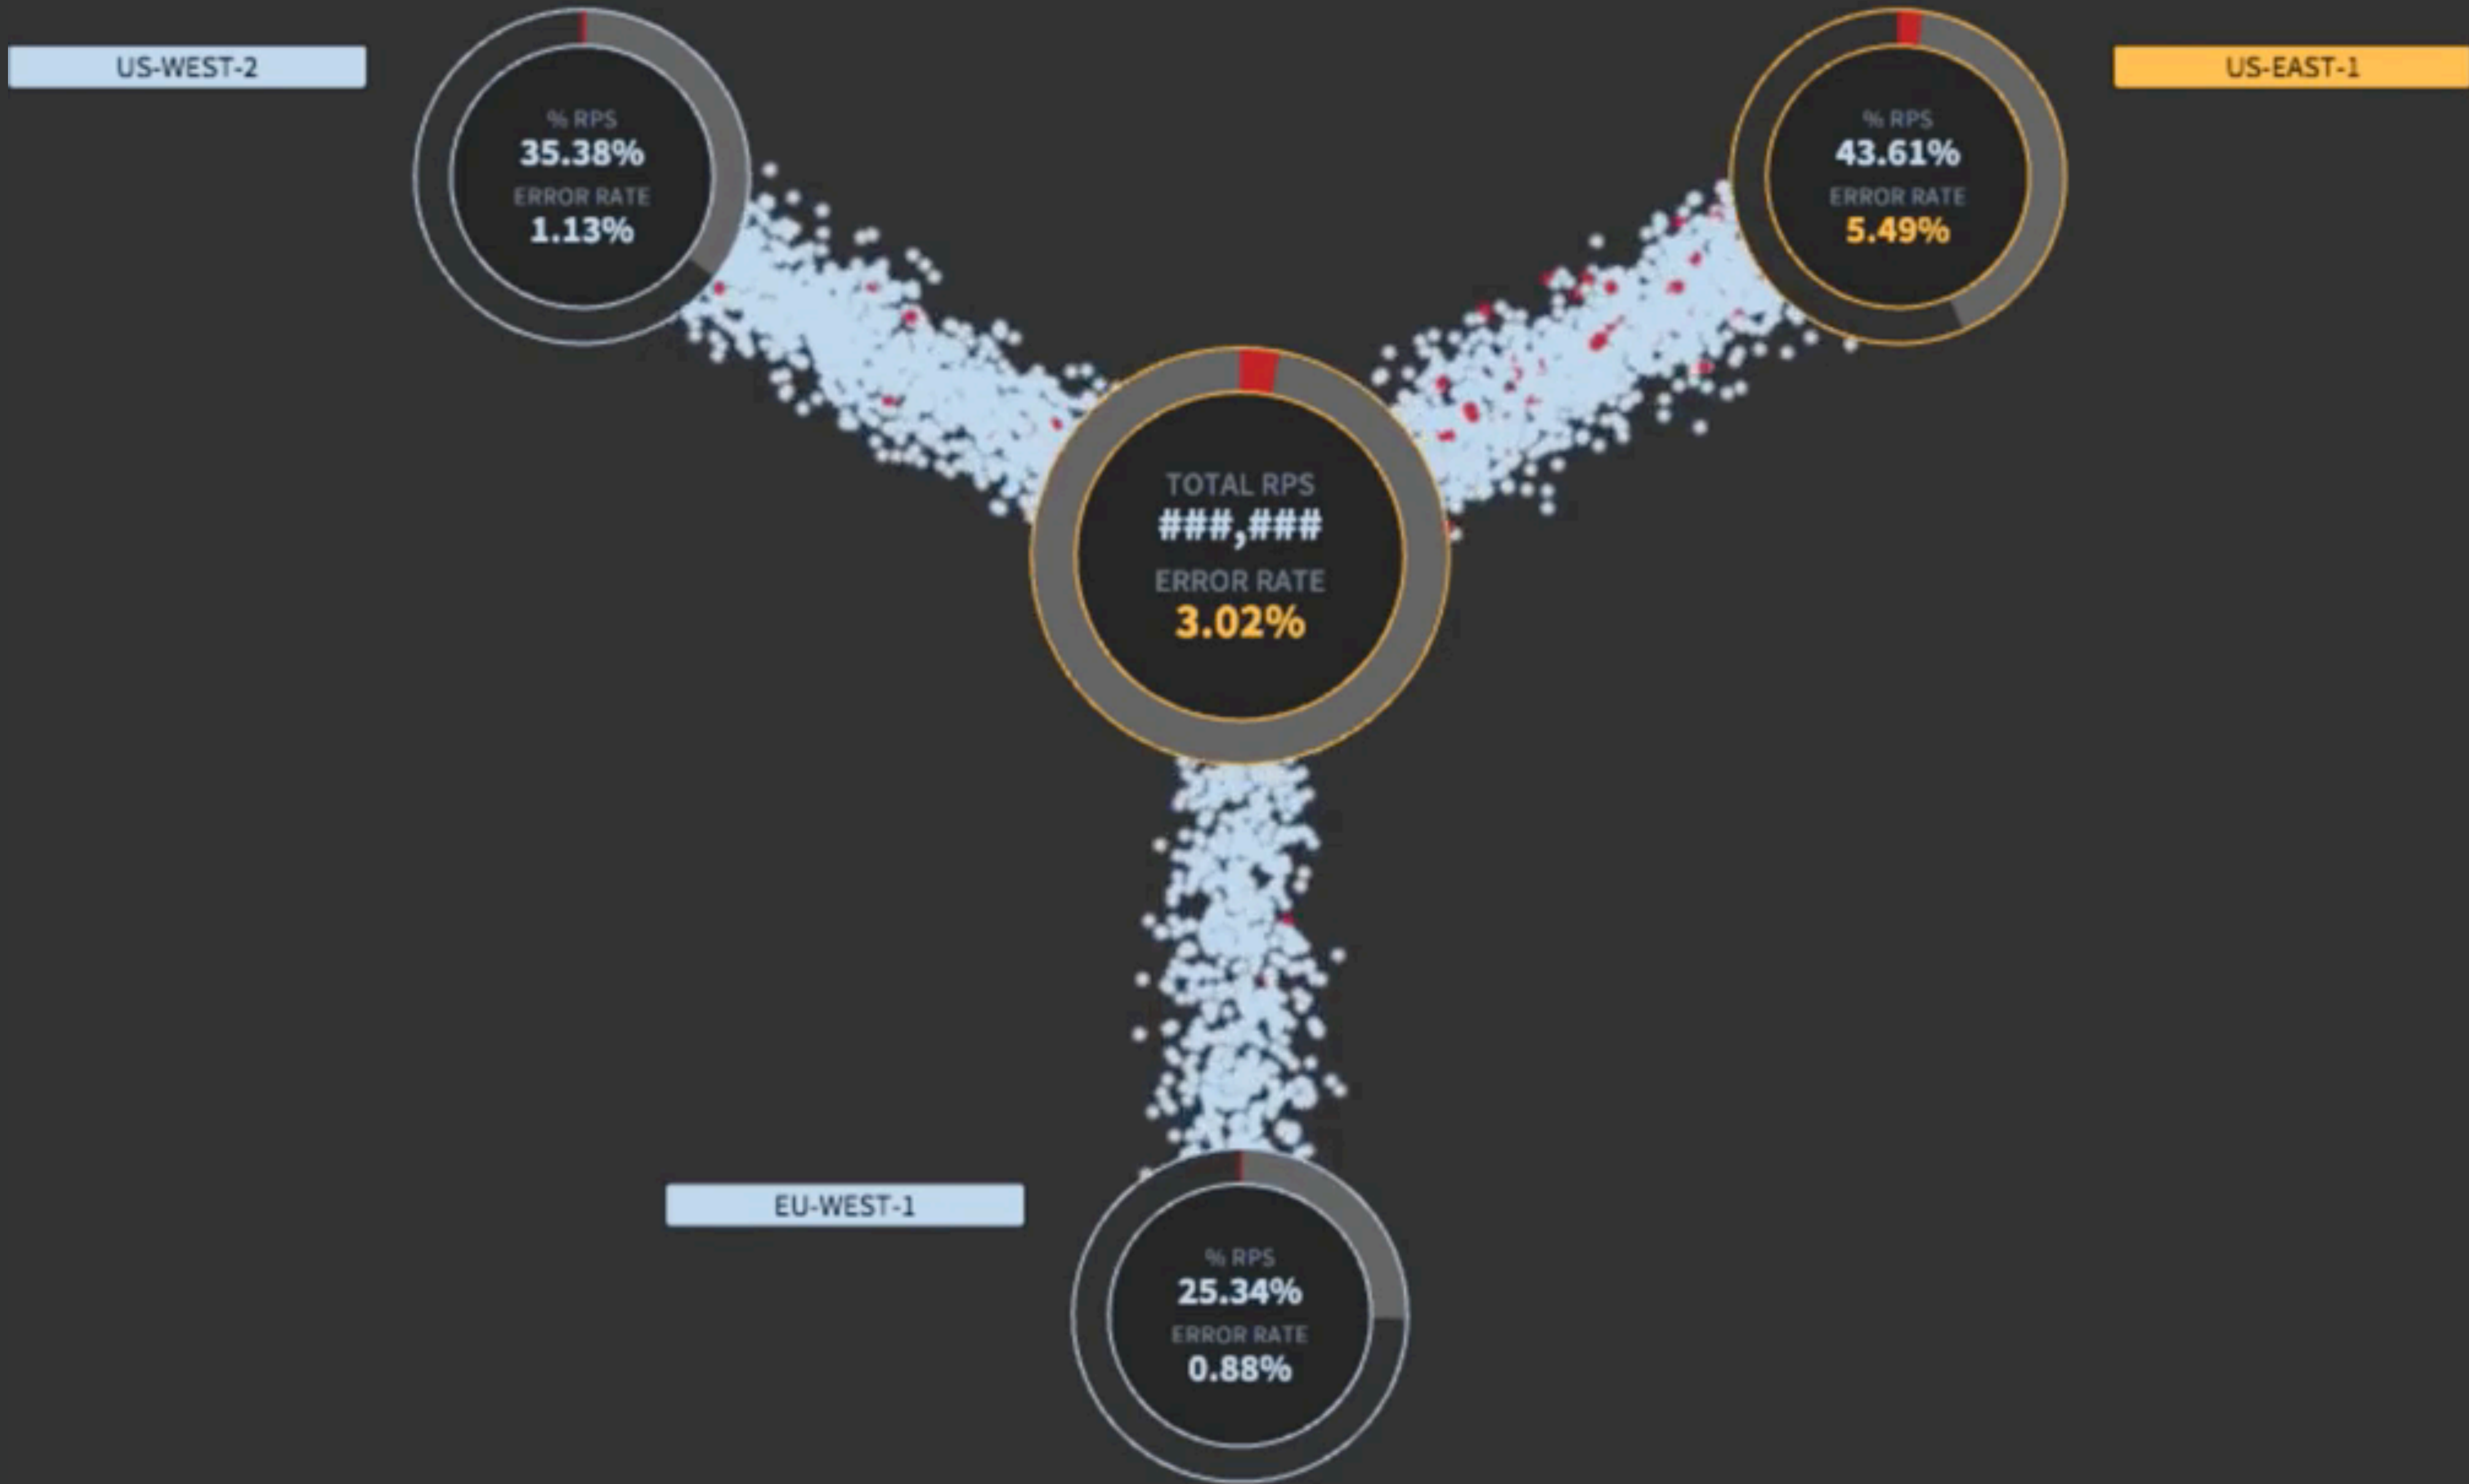


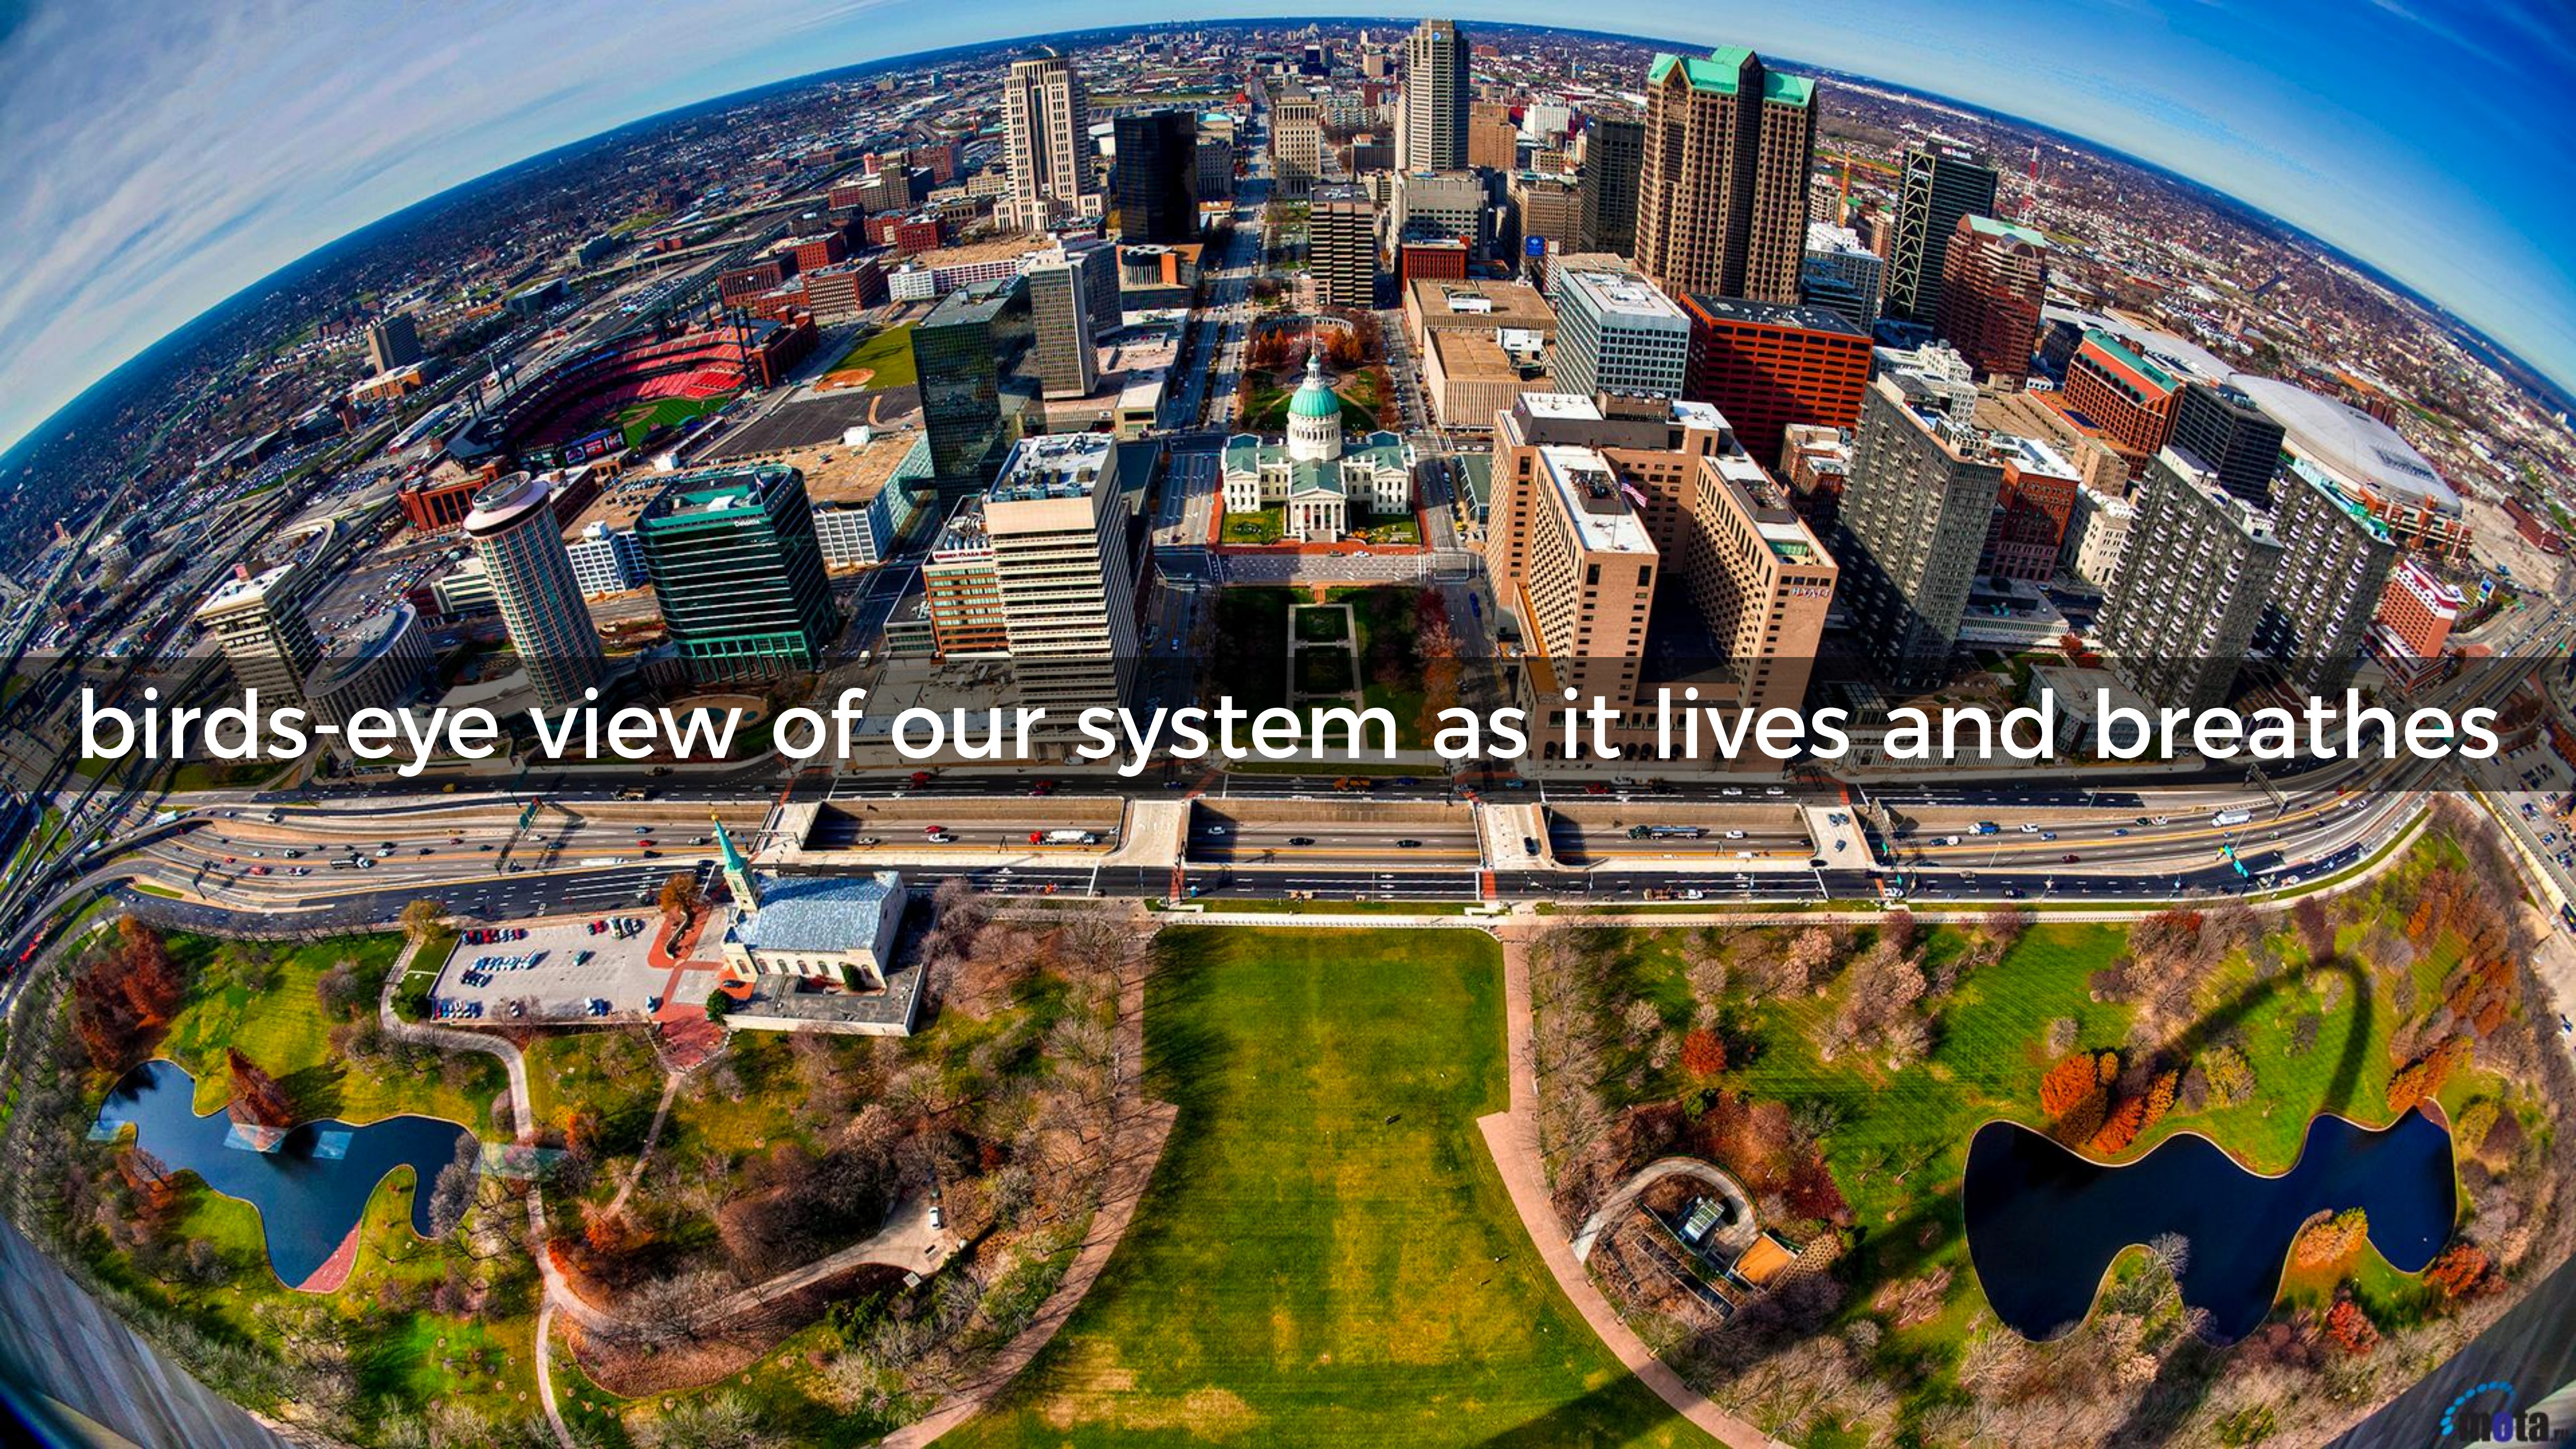




Service Traffic Map

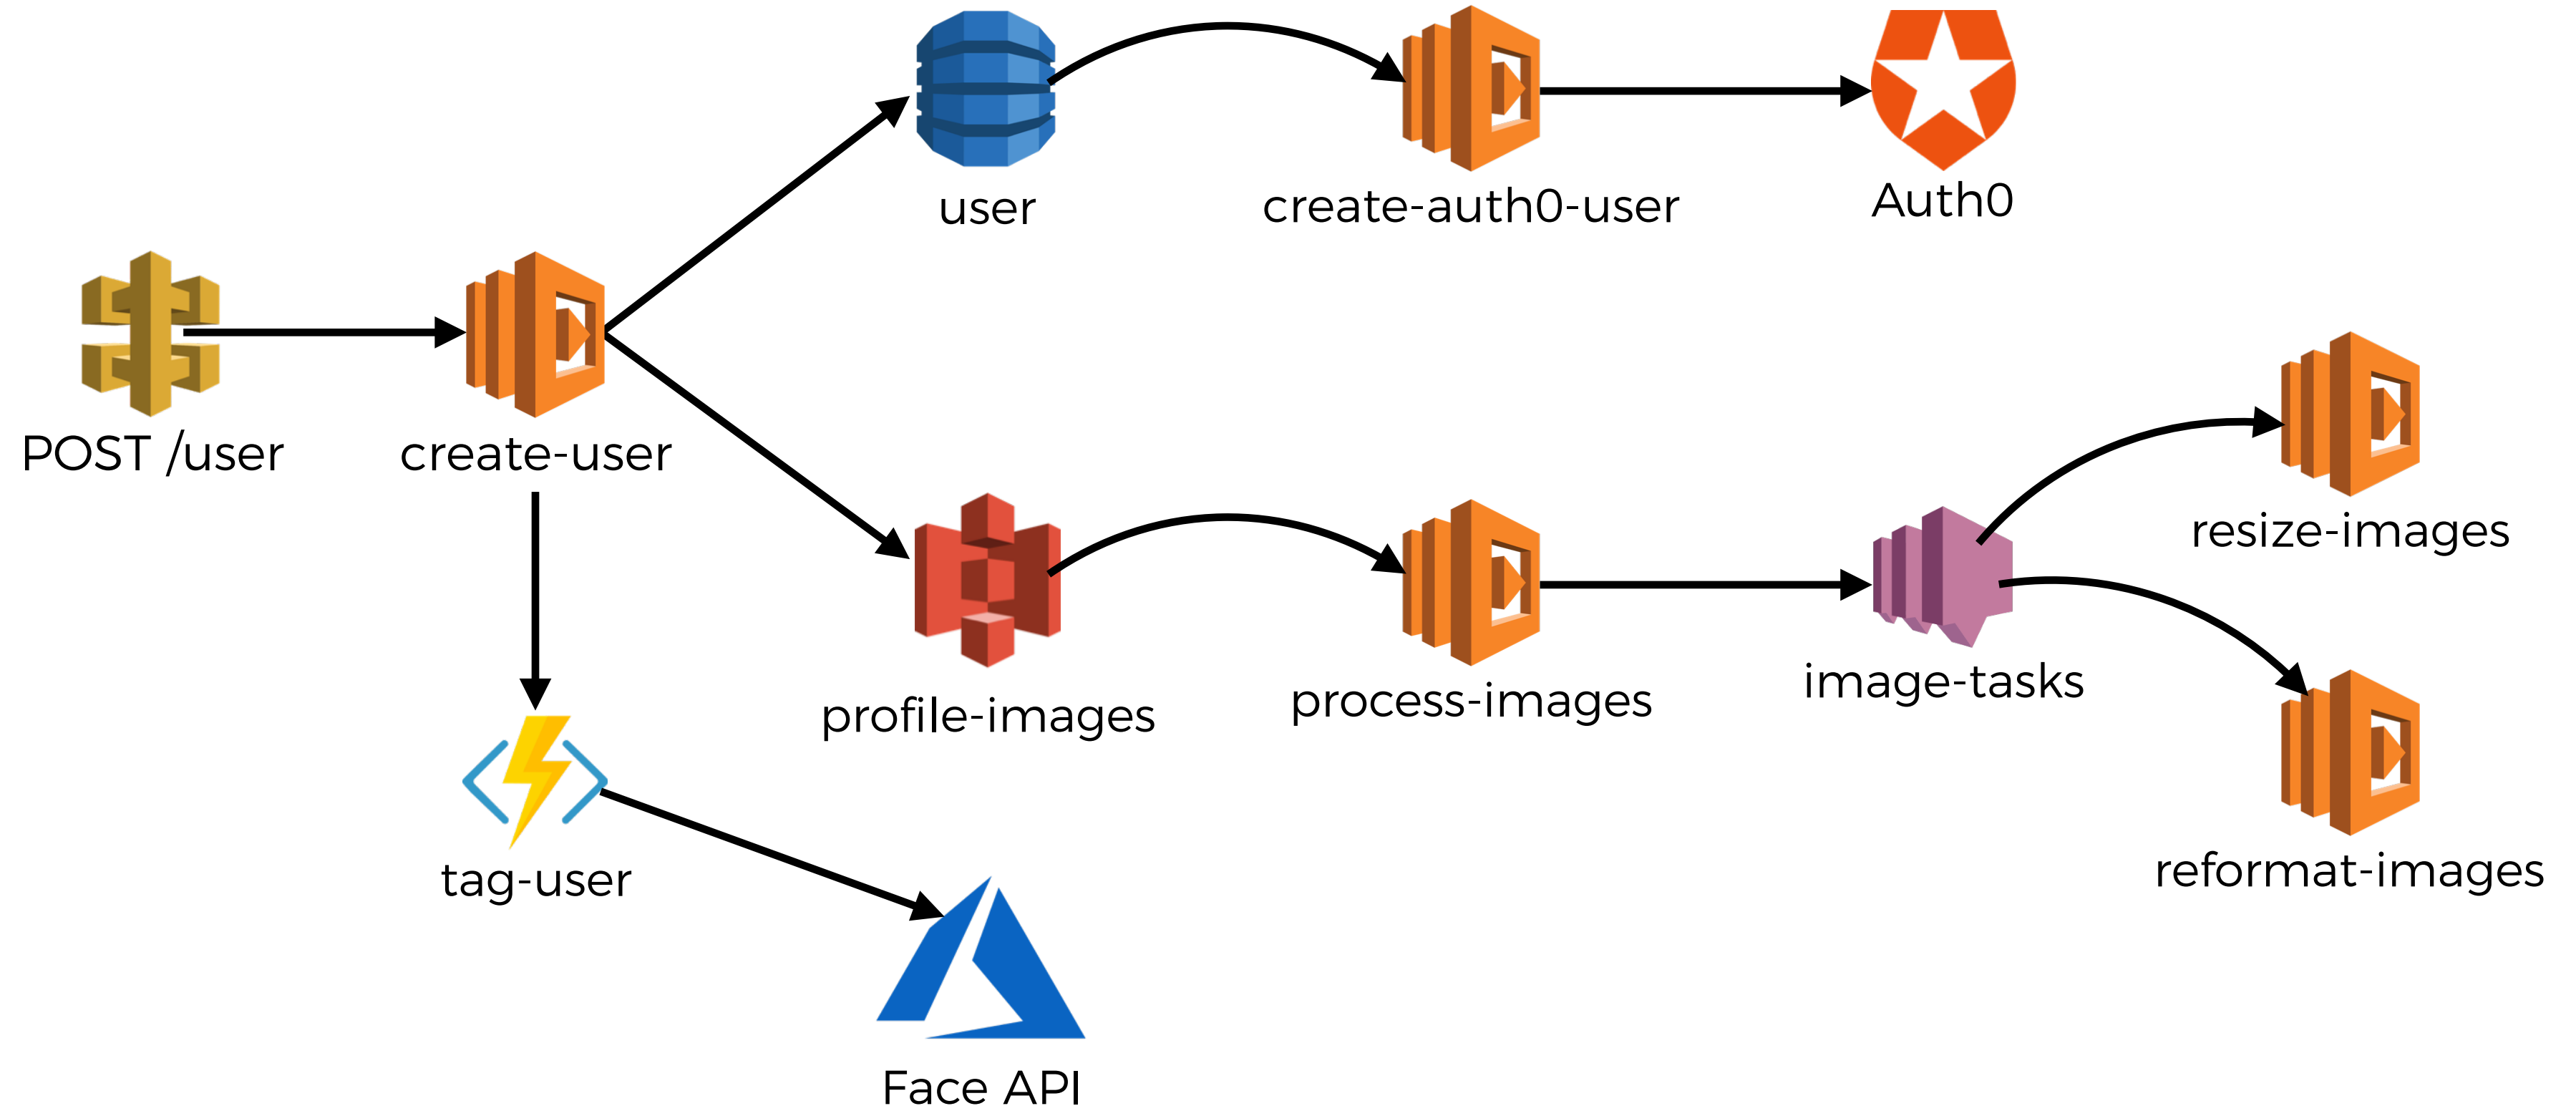
Filters ▾ Display ▾

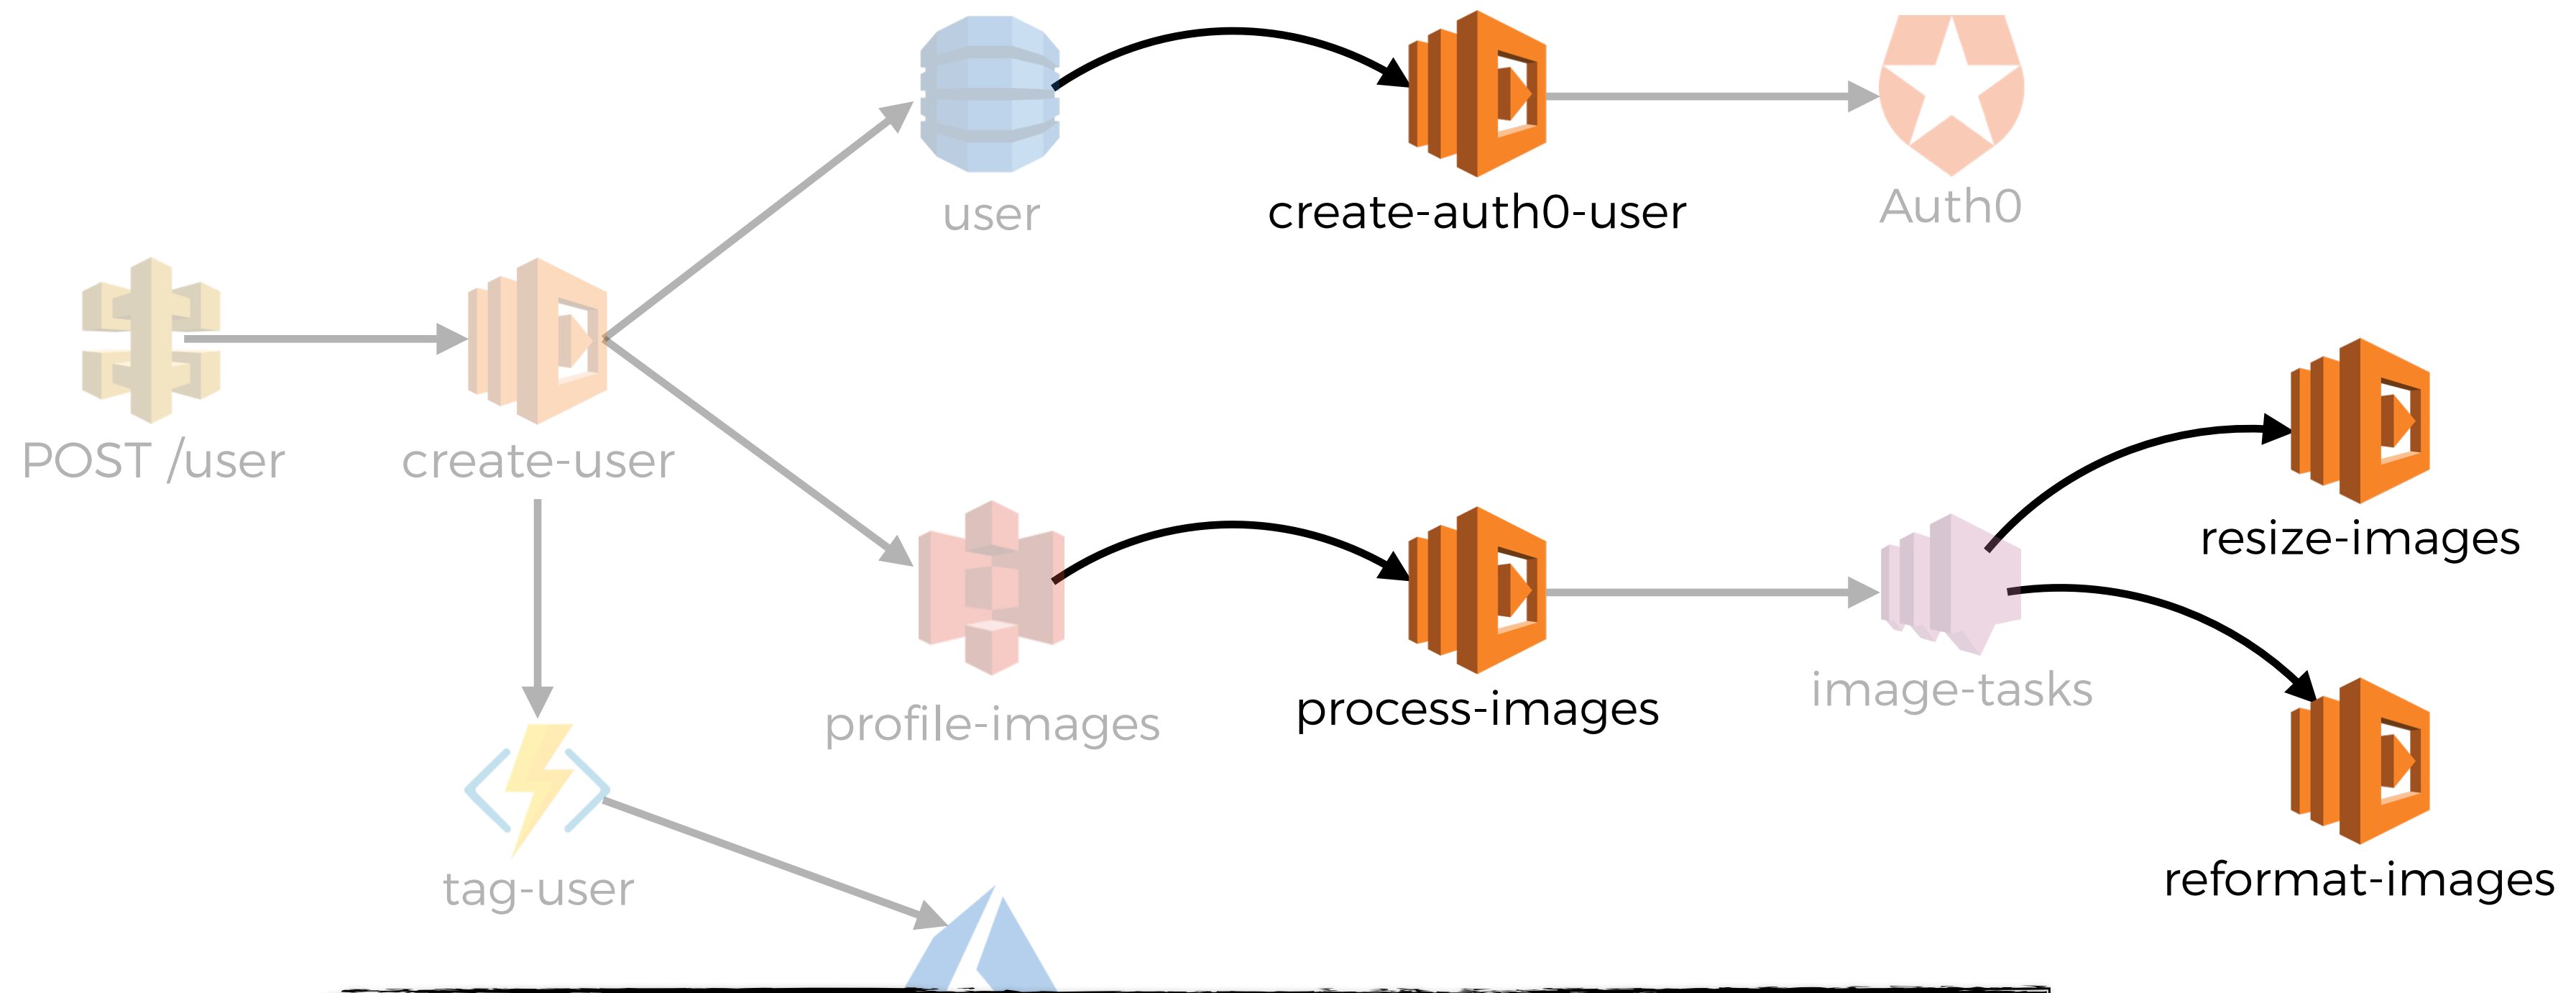




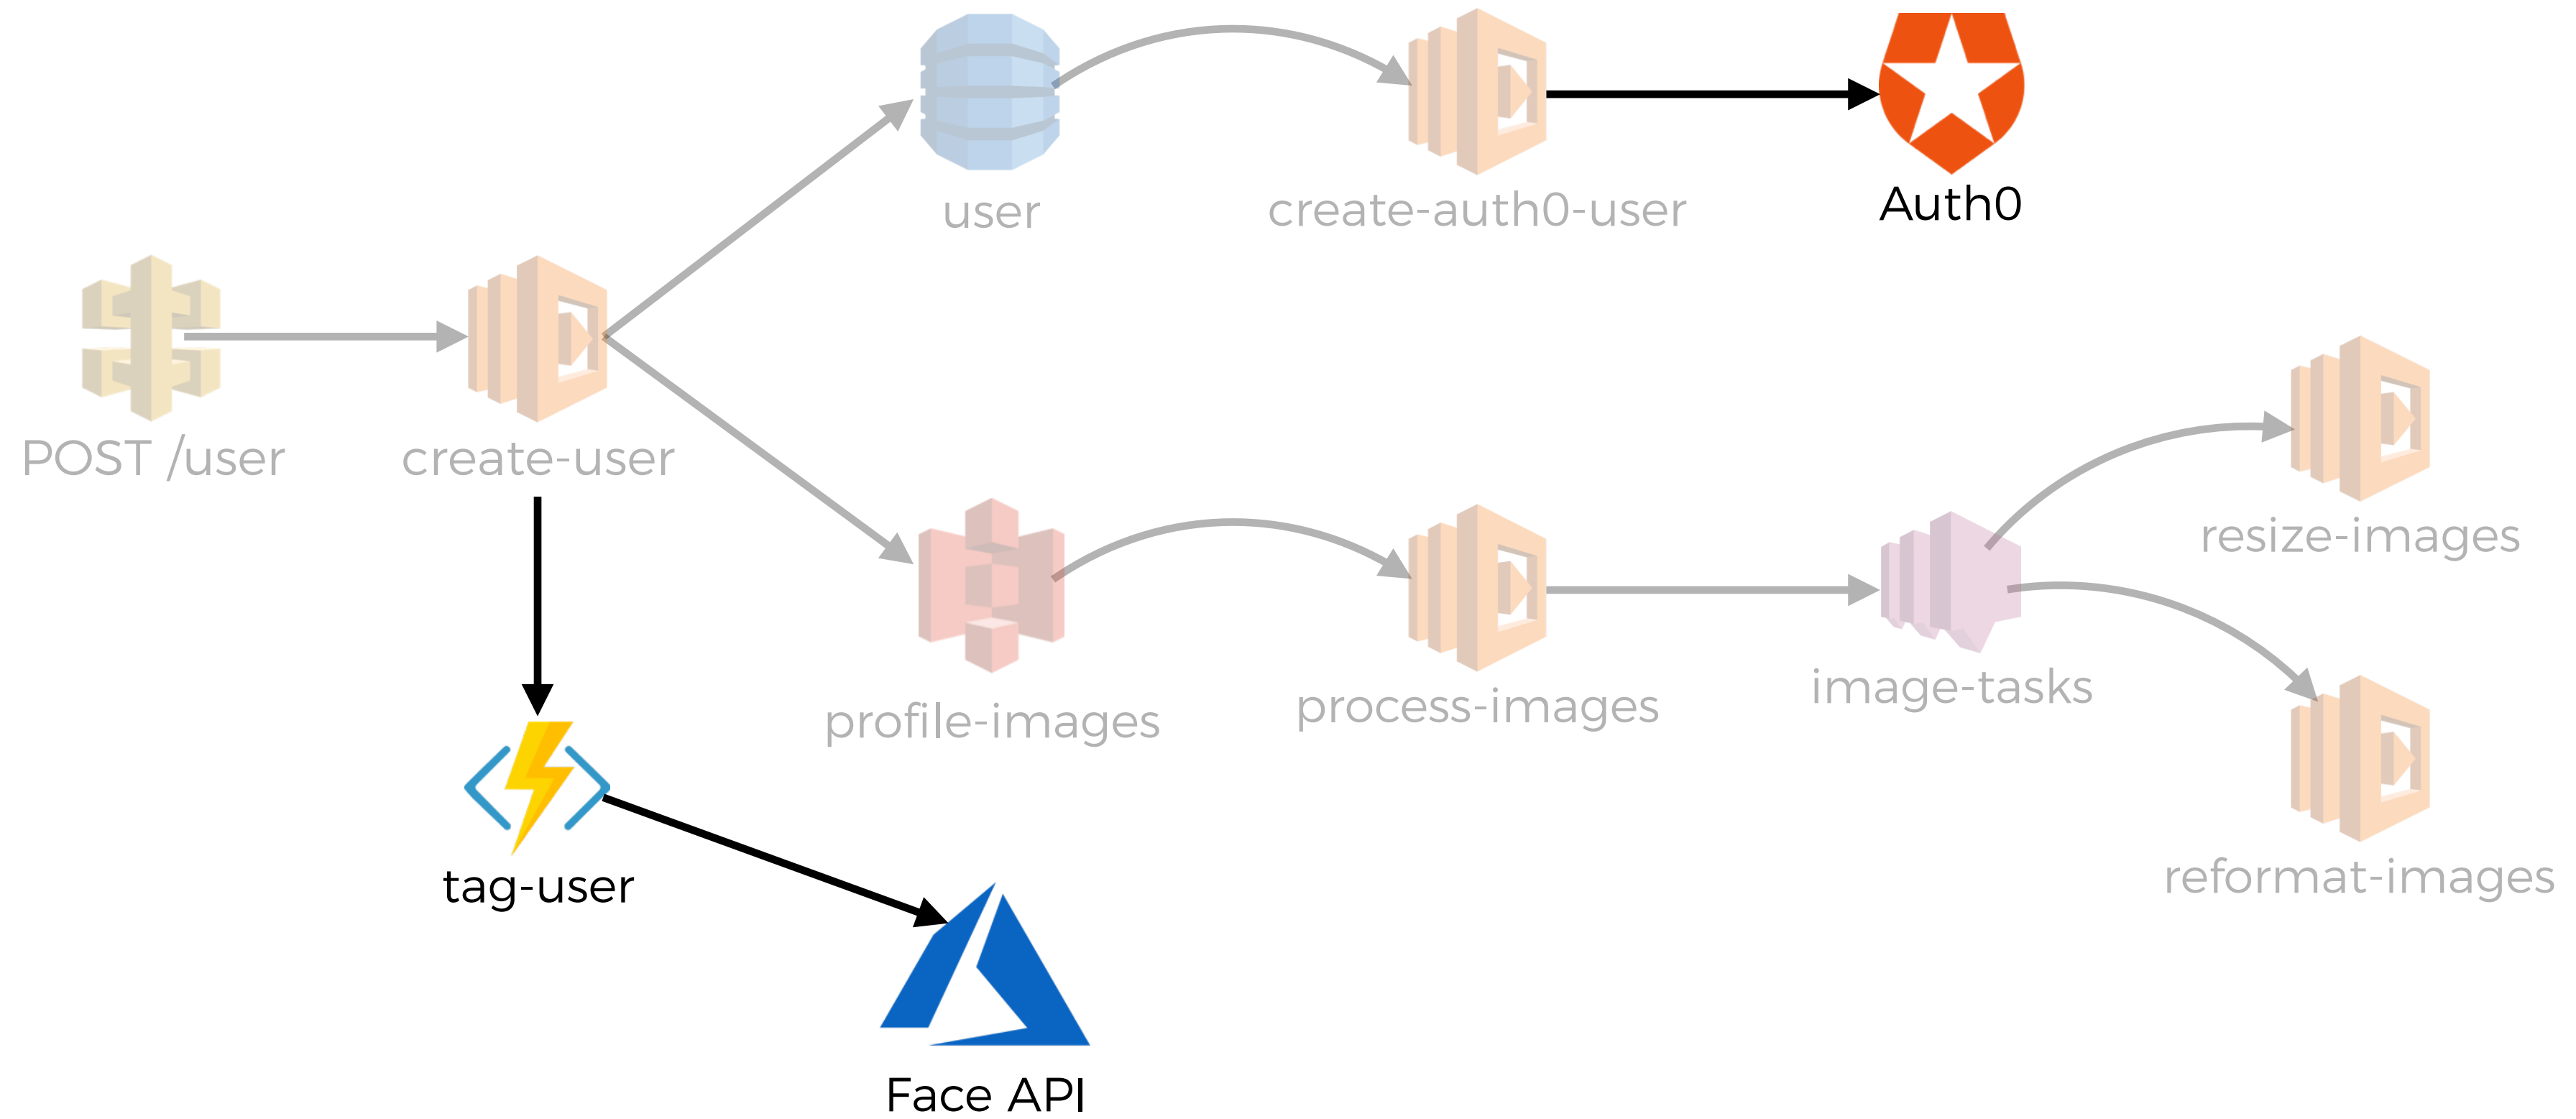
birds-eye view of our system as it lives and breathes



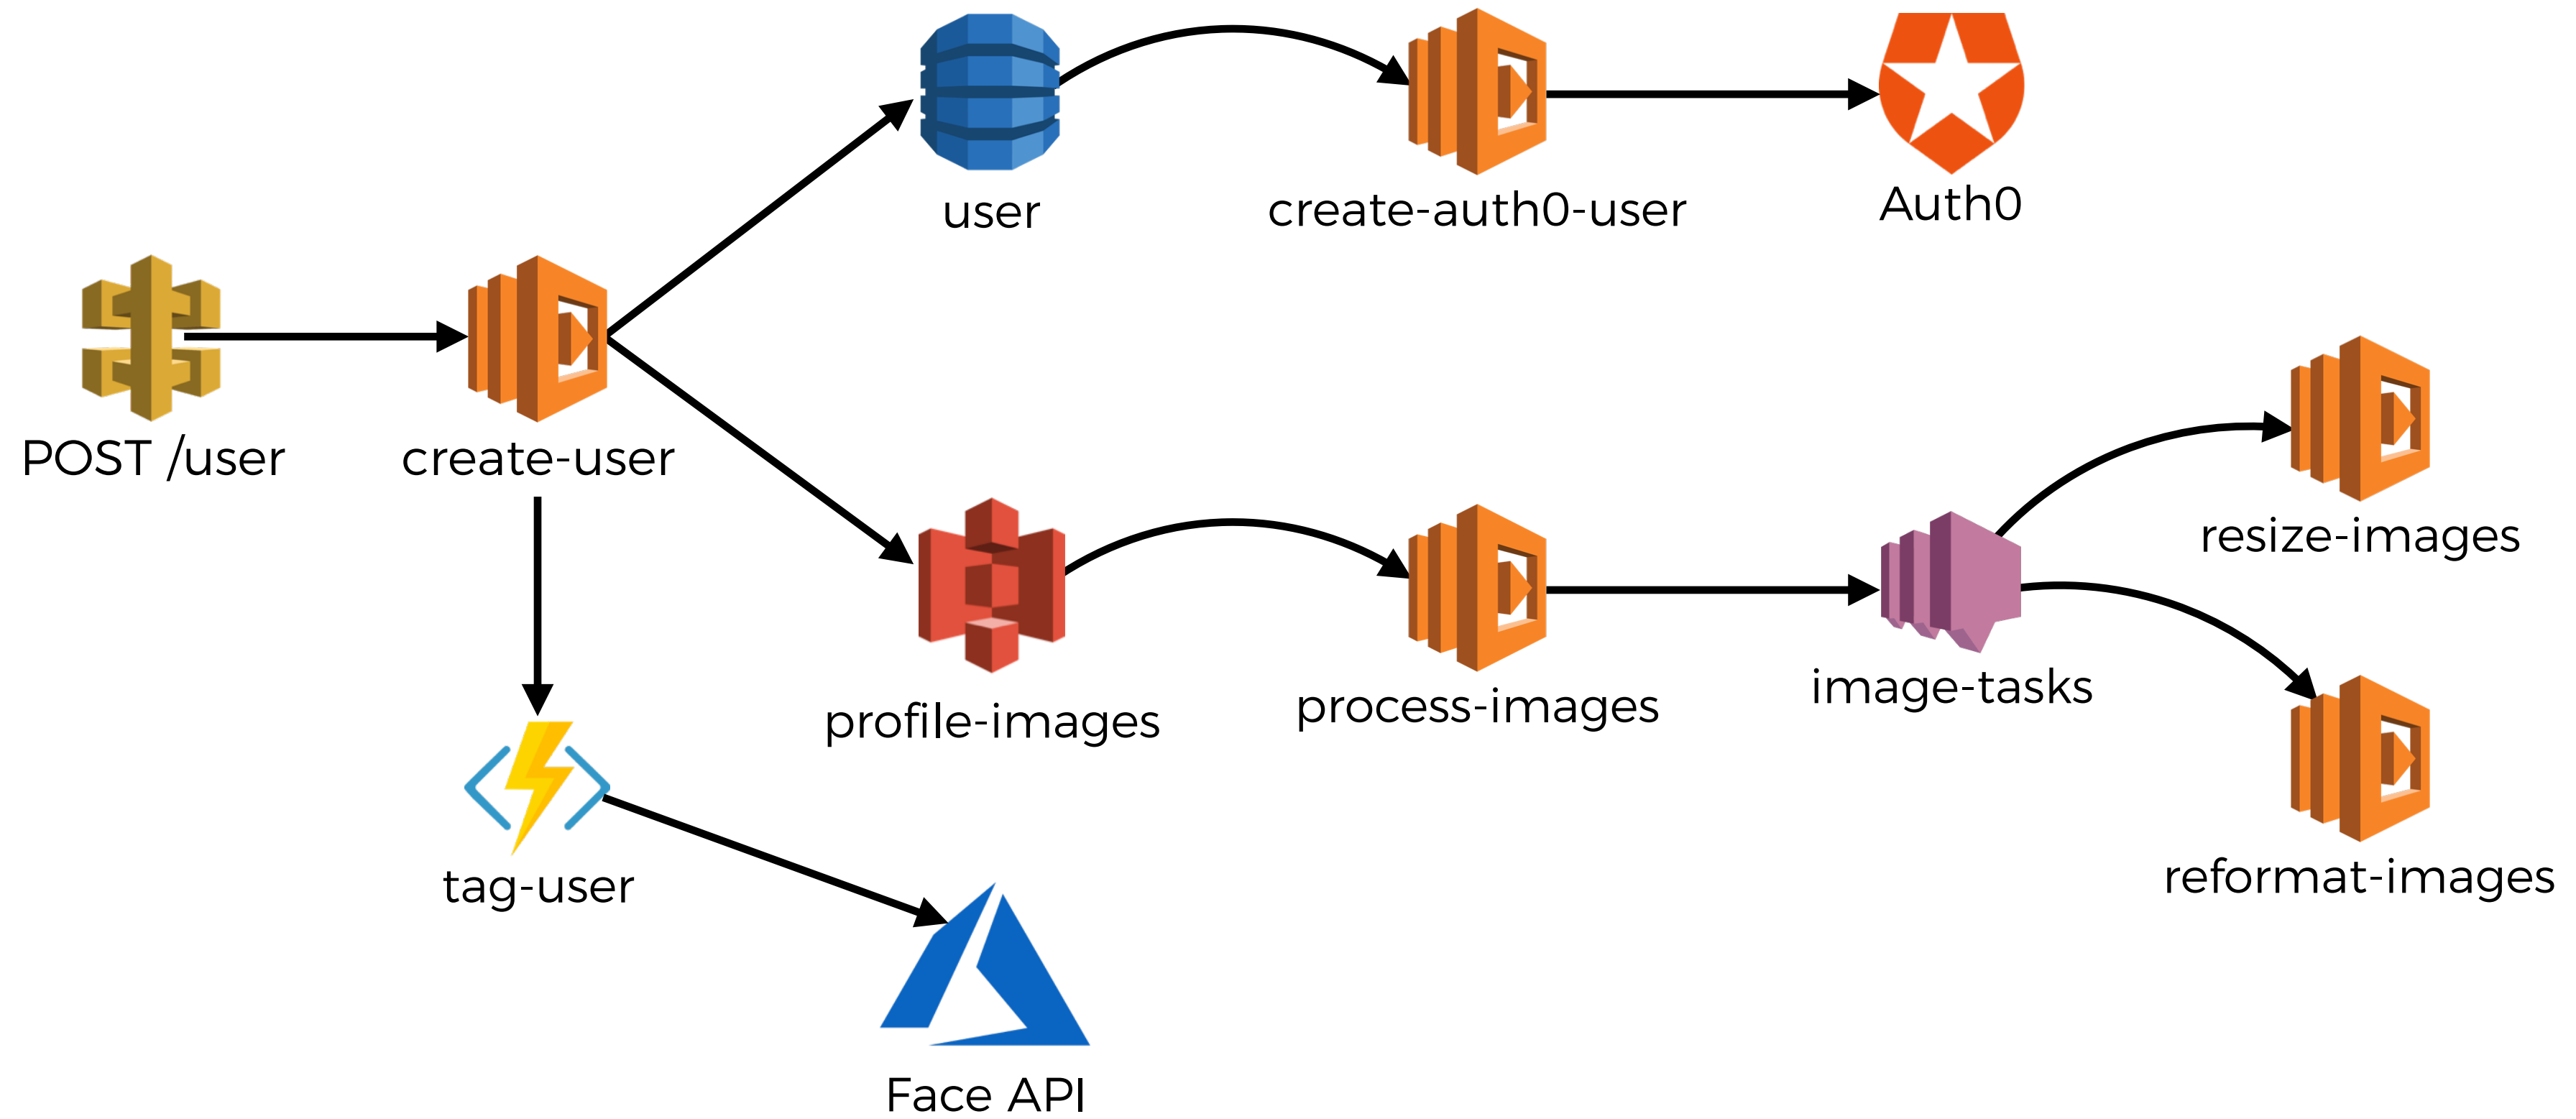




trace **async** invocations

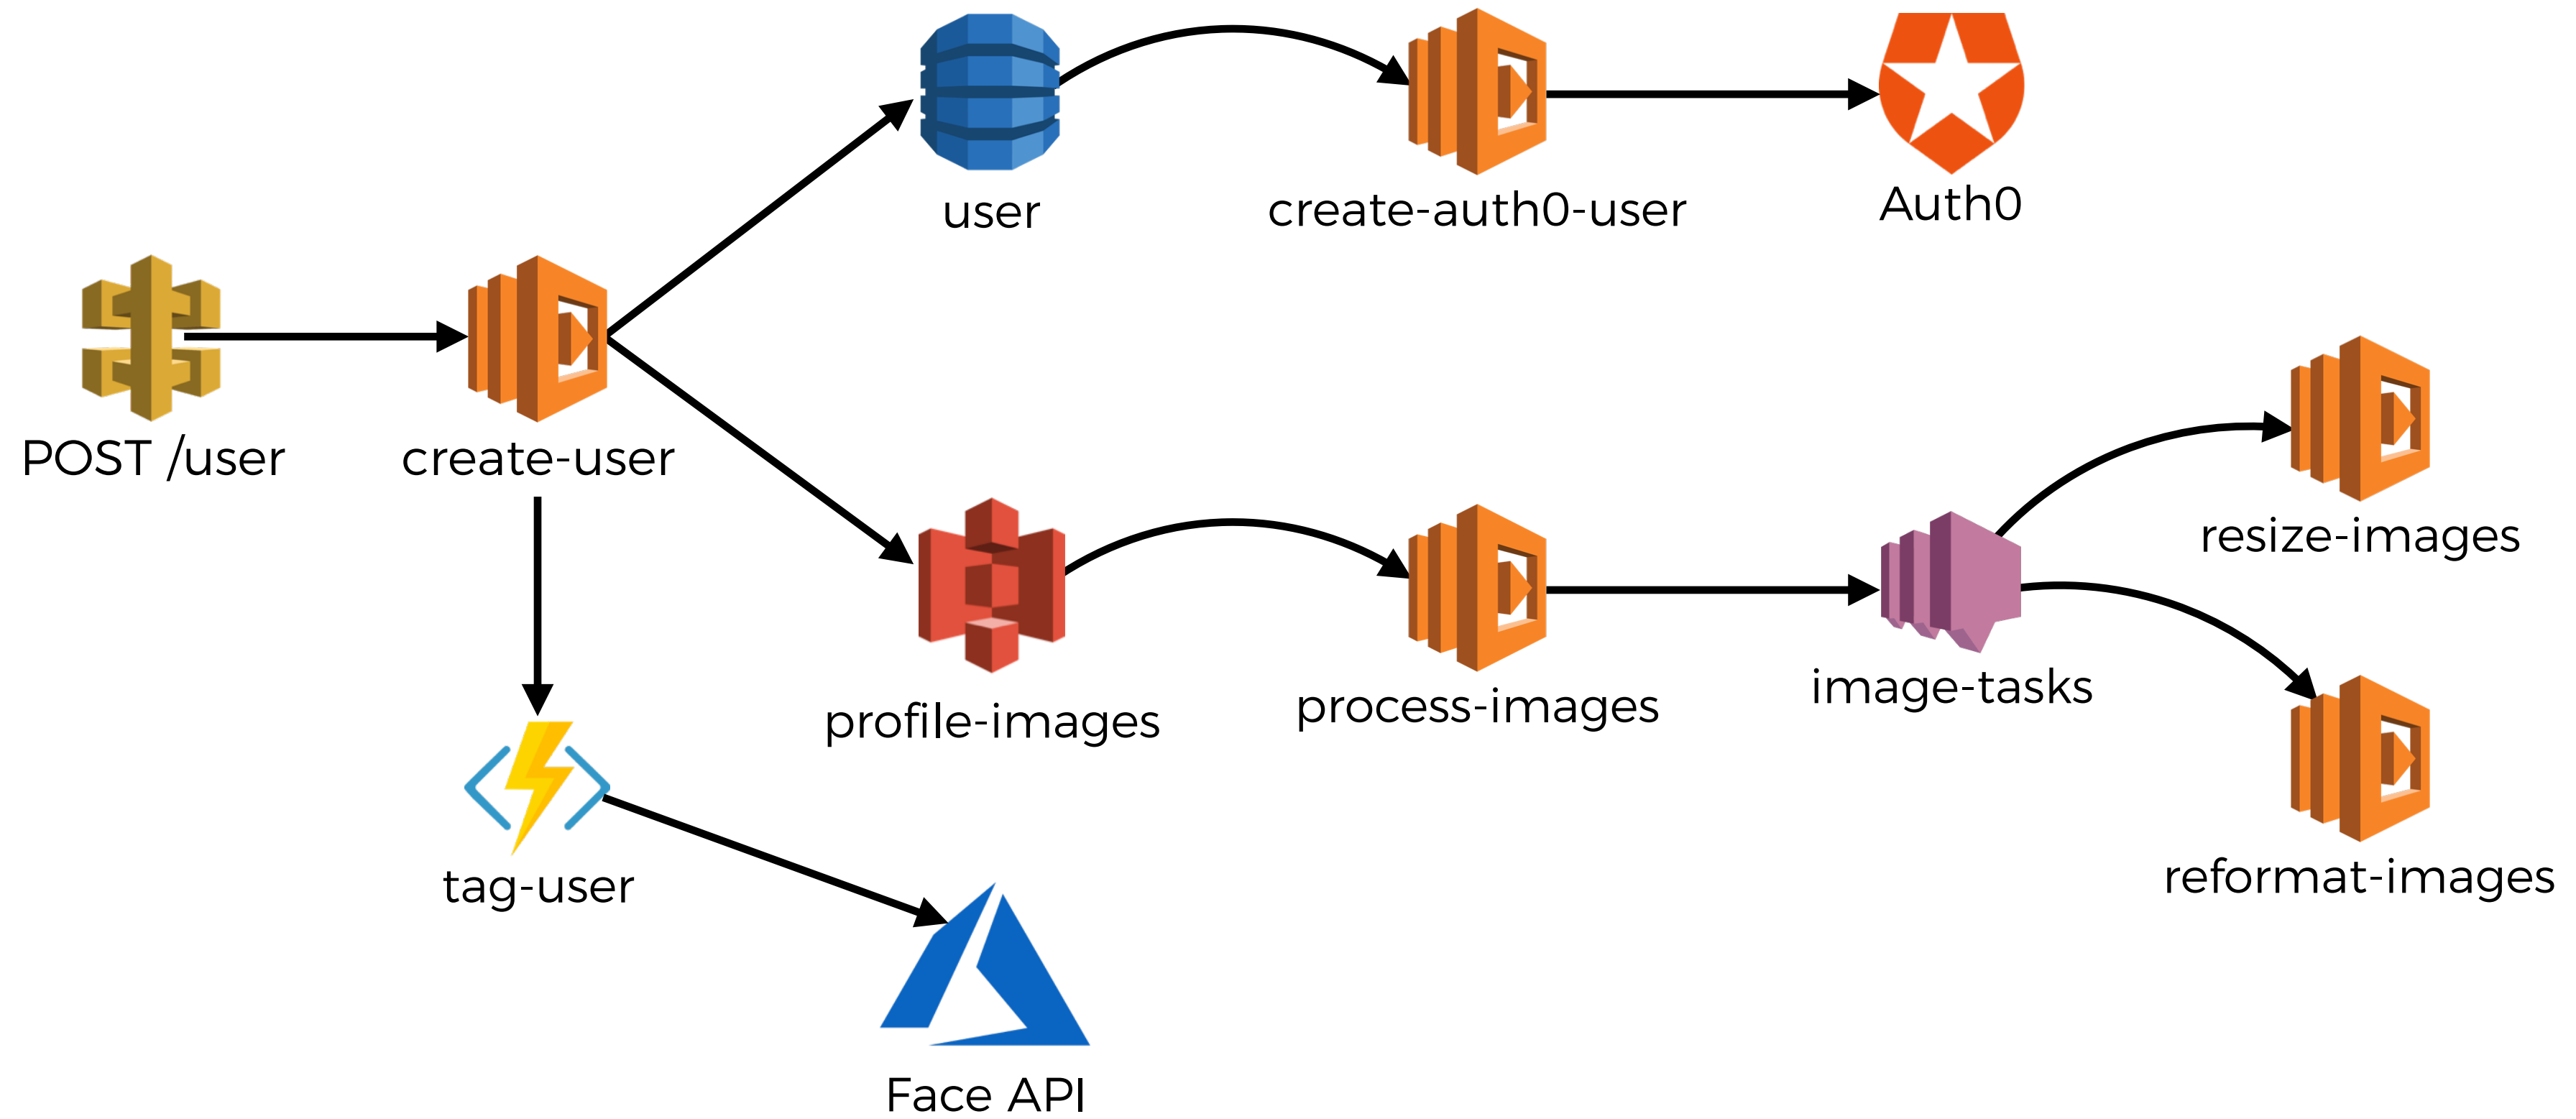


trace **non-AWS** resources



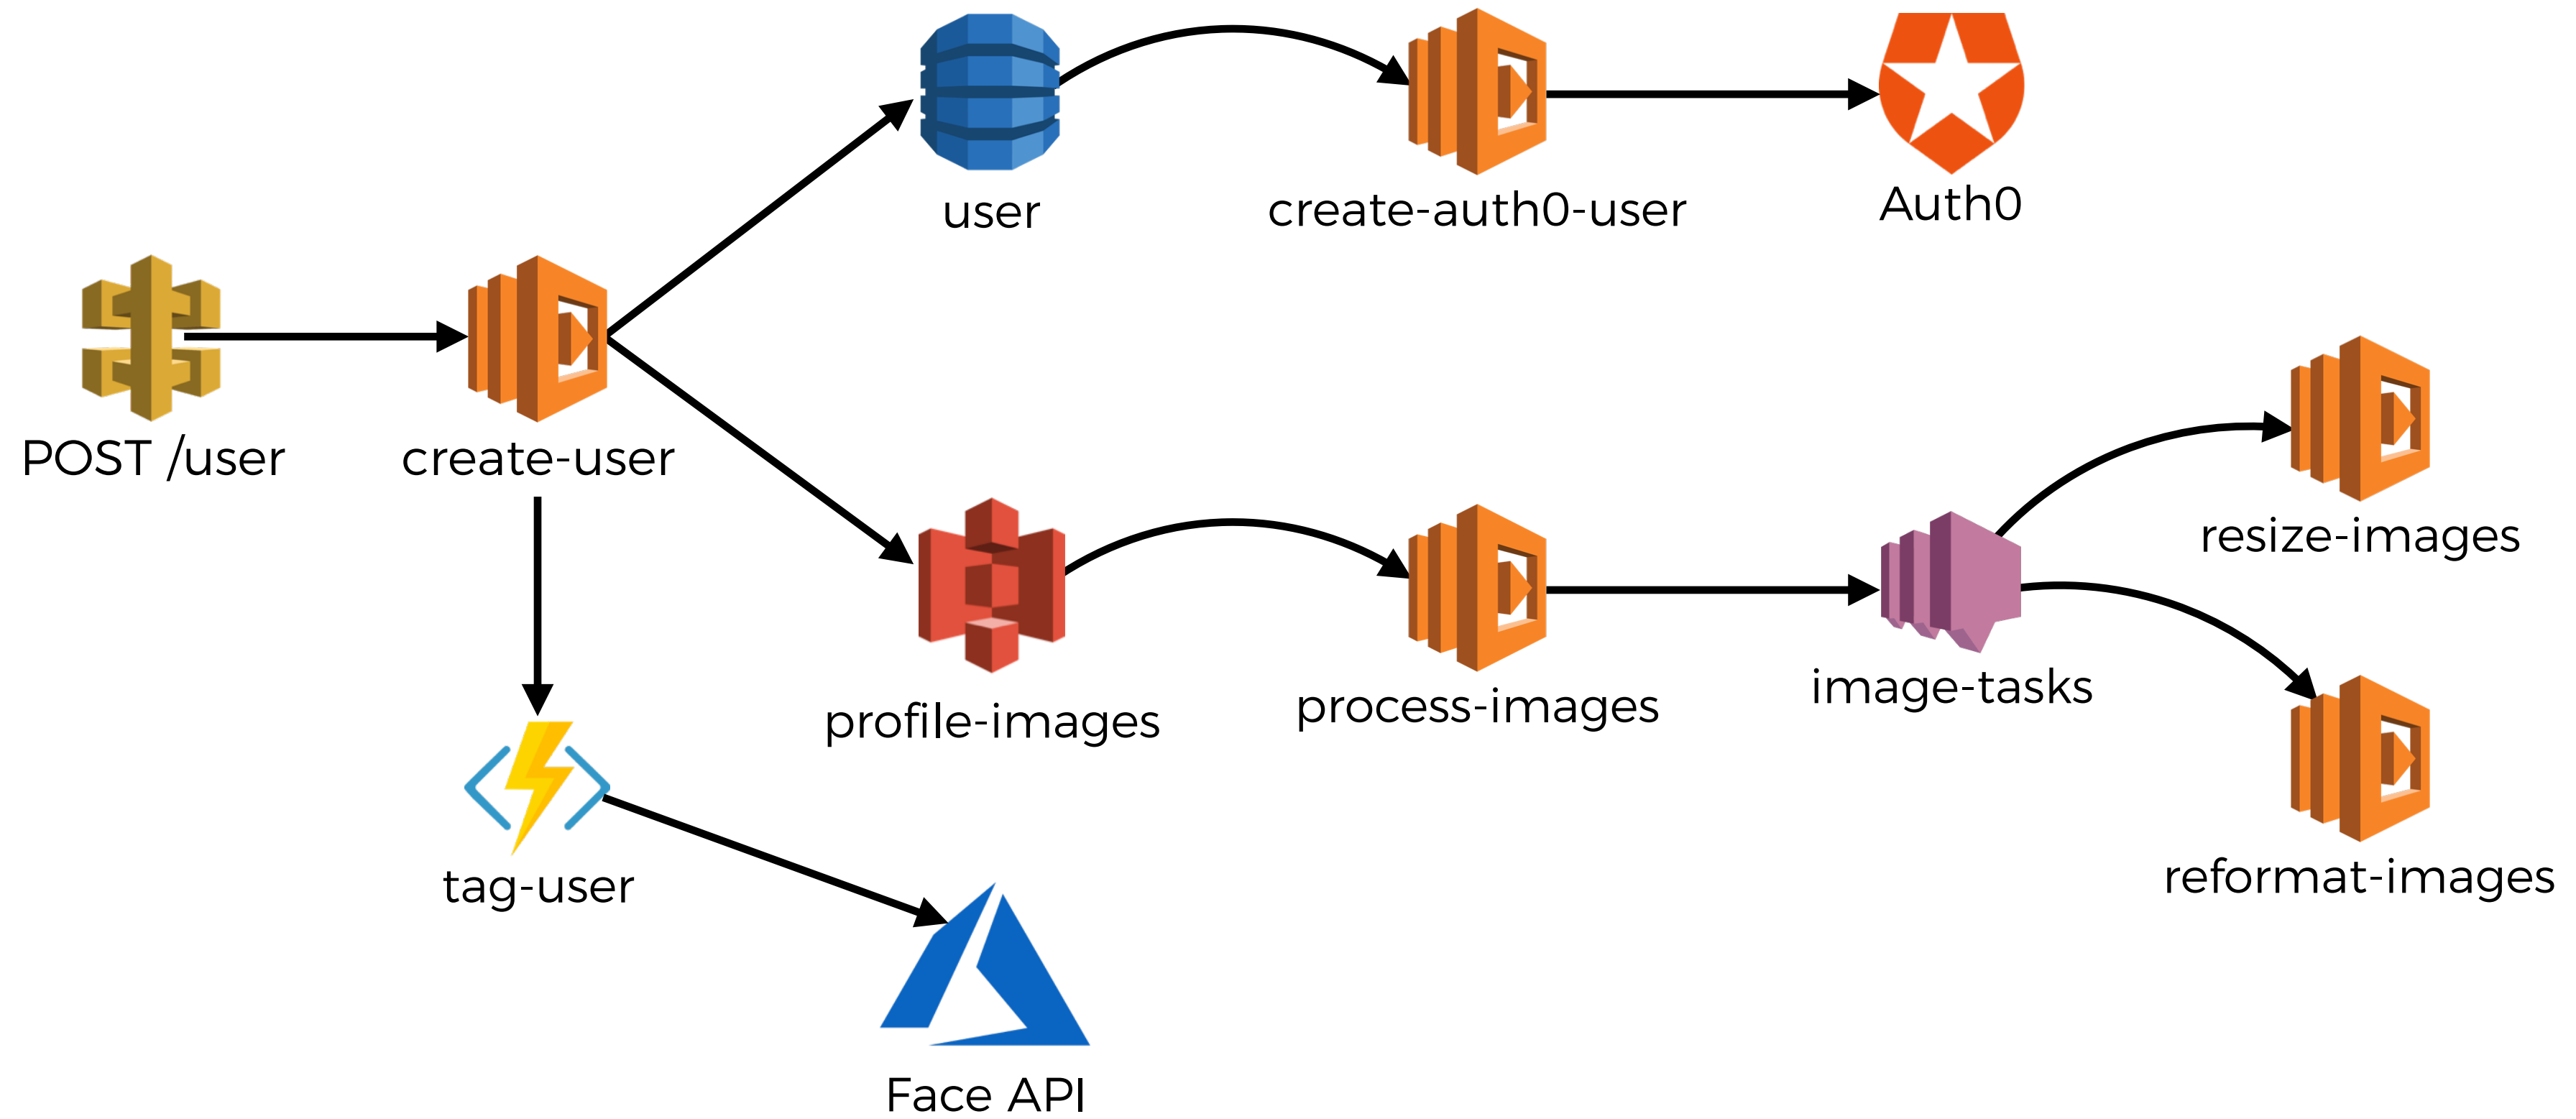
Logs

timestamp	component	level	message
2018/01/25 20:51:23.188	POST /user	debug	incoming request...
2018/01/25 20:51:23.201	create-user	debug	saving user [theburningmonk] in the [user] table...
2018/01/25 20:51:23.215	create-user	debug	saved user [theburningmonk] in the [user] table
2018/01/25 20:51:23.521	tag-user	debug	tagging user [theburningmonk] with Azure Face API...



Logs

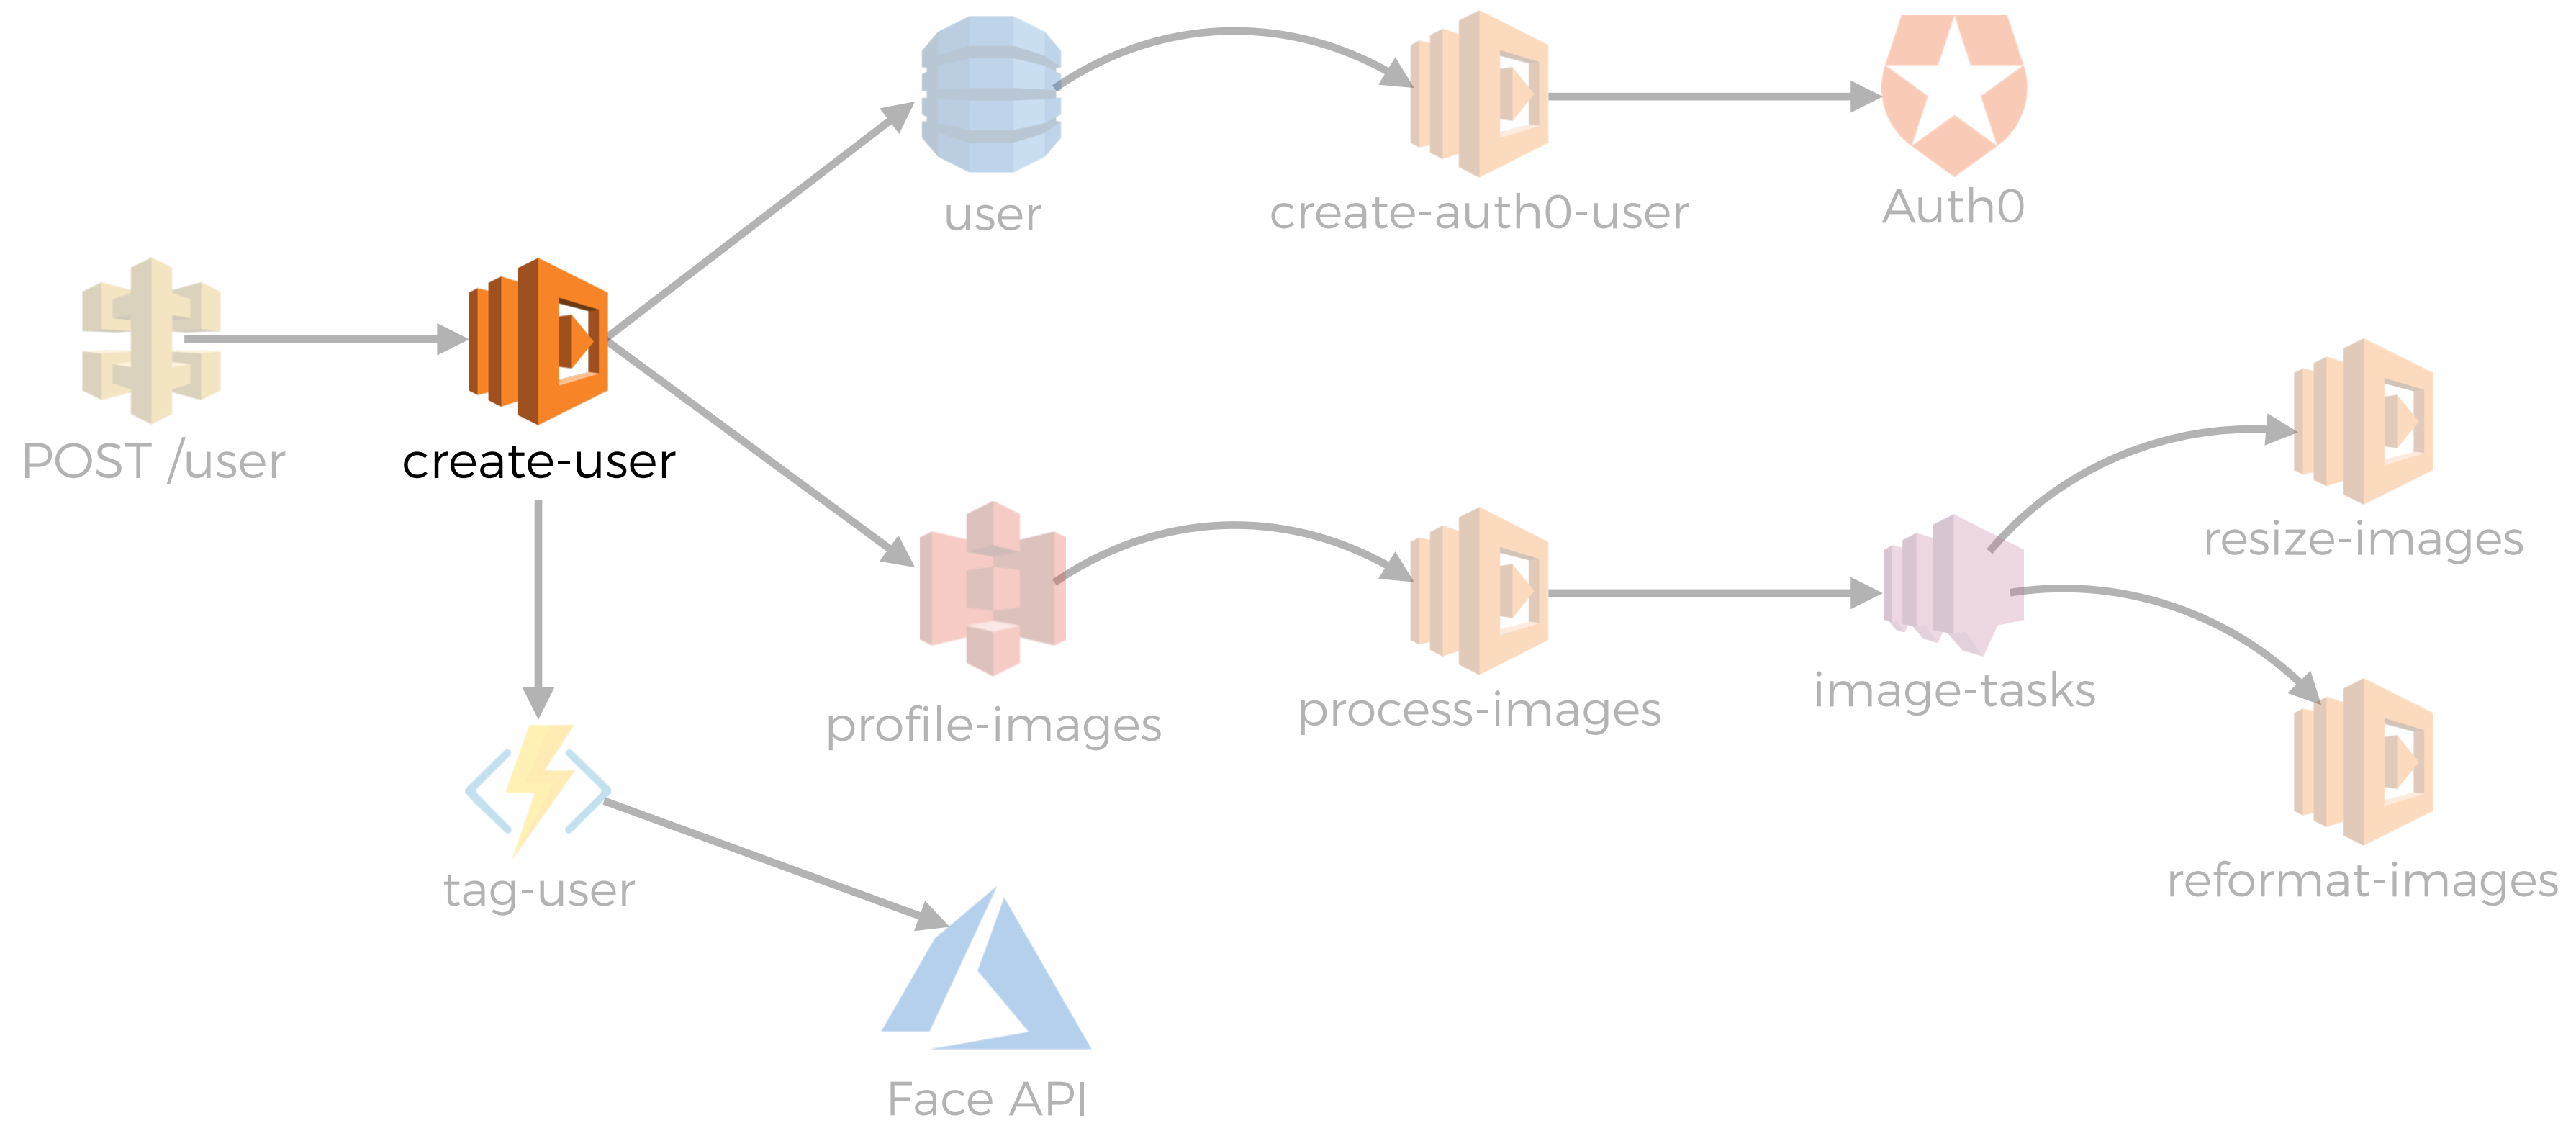
timestamp	component	level	message
2018/01/25 20:51:23.188	POST /user	debug	incoming request...
2018/01/25 20:51:23.201	create-user	debug	saving user [theburningmonk] in the [user] table...
2018/01/25 20:51:23.215	create-user	debug	saved user [theburningmonk] in the [user] table
2018/01/25 20:51:23.521	tag-user	debug	tagging user [theburningmonk] with Azure Face API...



Logs

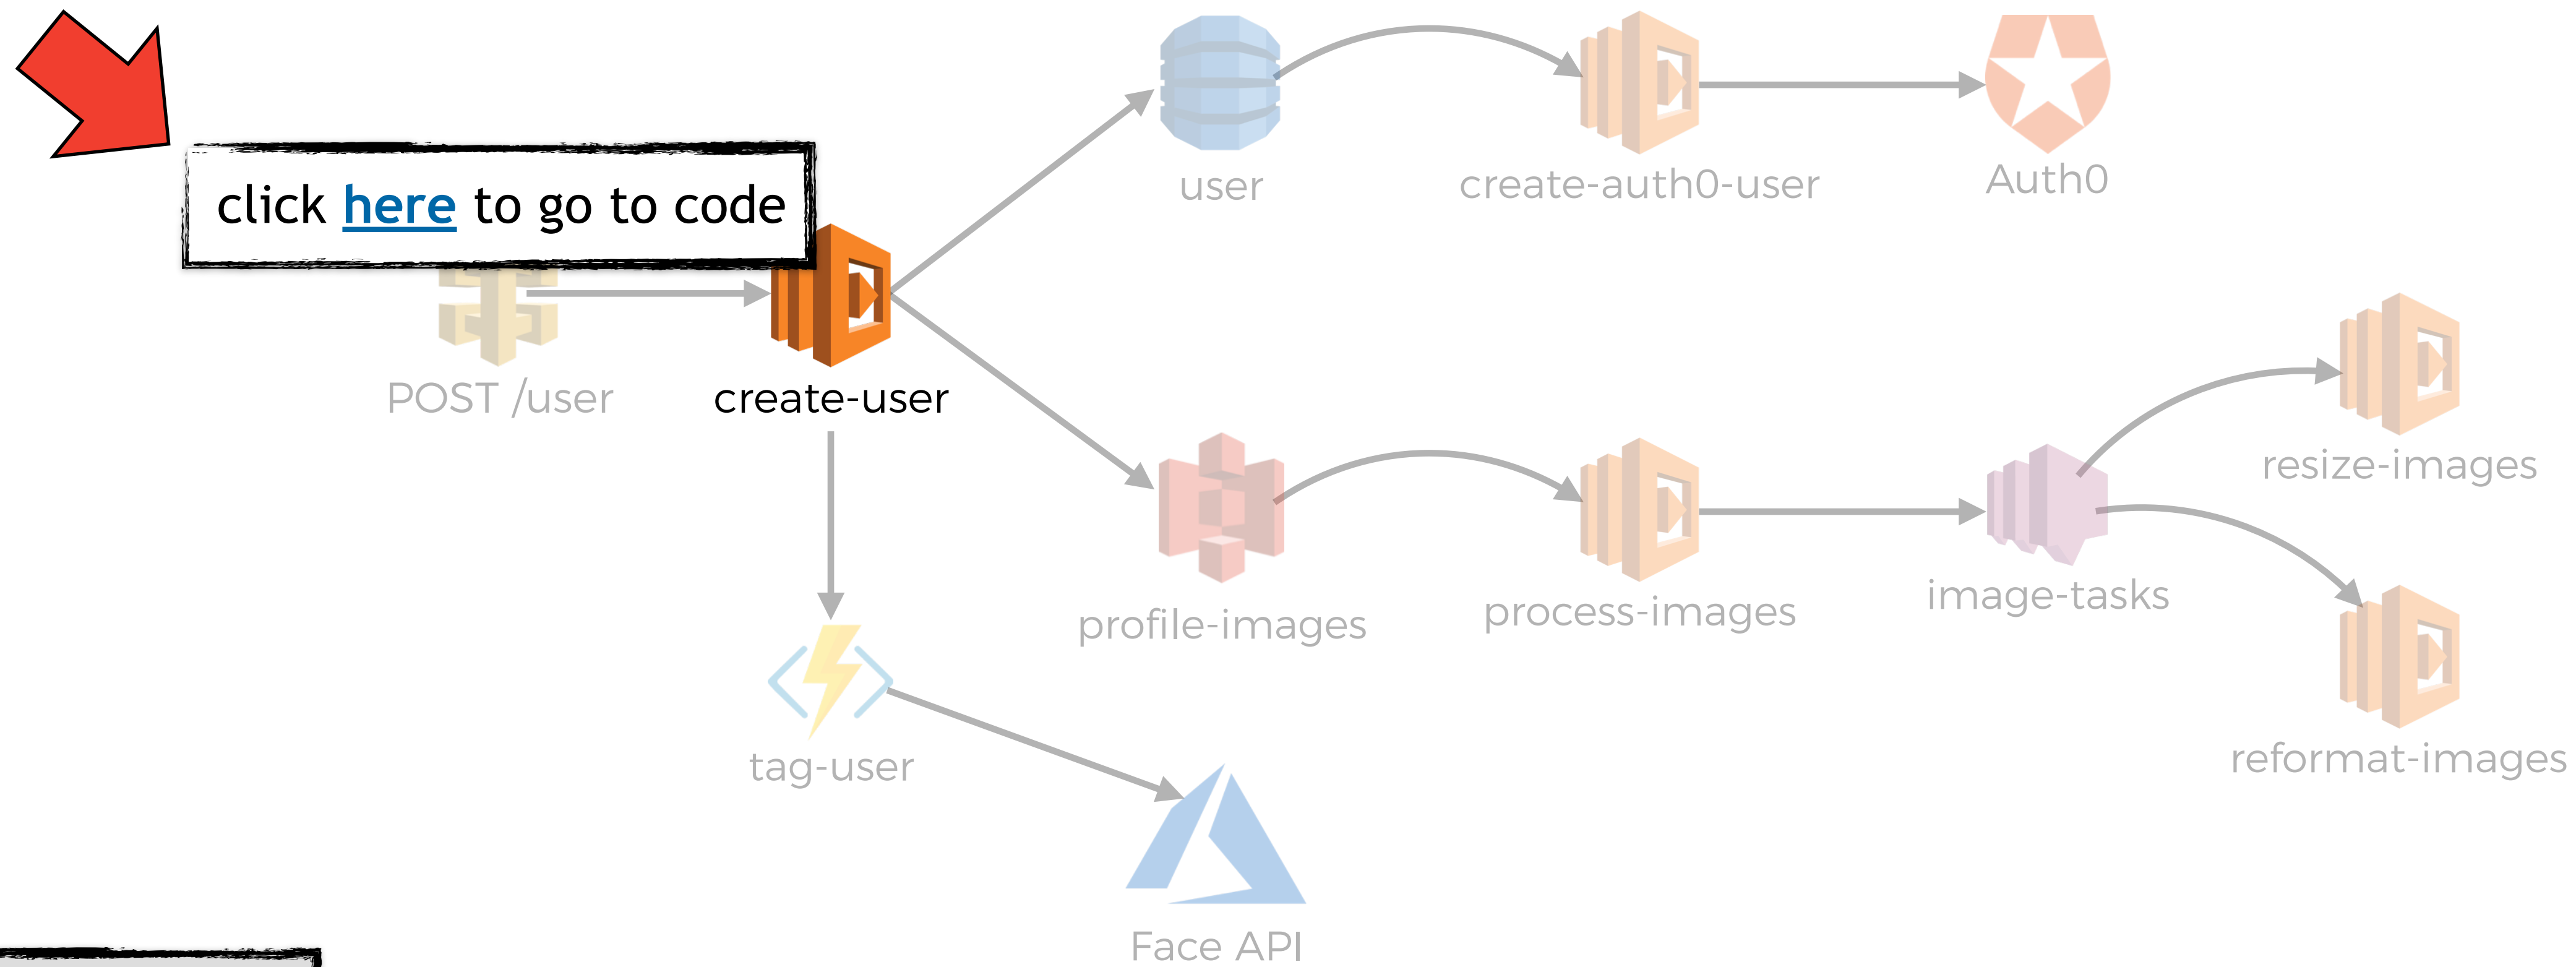
timestamp	component	level	message
2018/01/25 20:51:23.188	POST /user	debug	incoming request...

request-id	0ae4ba5d-dab1-4f9e-9de7-eace27ebfbc2
start-time	2018/01/25 20:51:23.188
method	POST



Logs

timestamp	component	level	message
2018/01/25 20:51:23.201	create-user	debug	saving user [theburningmonk] in the [user] table...
2018/01/25 20:51:23.215	create-user	debug	saved user [theburningmonk] in the [user] table
2018/01/25 20:51:23.585	create-user	debug	tagged user [theburningmonk] with Azure Face API...
2018/01/25 20:51:23.587	create-user	debug	uploading profile image...



click [here](#) to go to code

Logs

timestamp	component	level	message
2018/01/25 20:51:23.201	create-user	debug	saving user [theburningmonk] in the [user] table...
2018/01/25 20:51:23.215	create-user	debug	saved user [theburningmonk] in the [user] table
2018/01/25 20:51:23.585	create-user	debug	tagged user [theburningmonk] with Azure Face API...
2018/01/25 20:51:23.587	create-user	debug	uploading profile image...



Environment

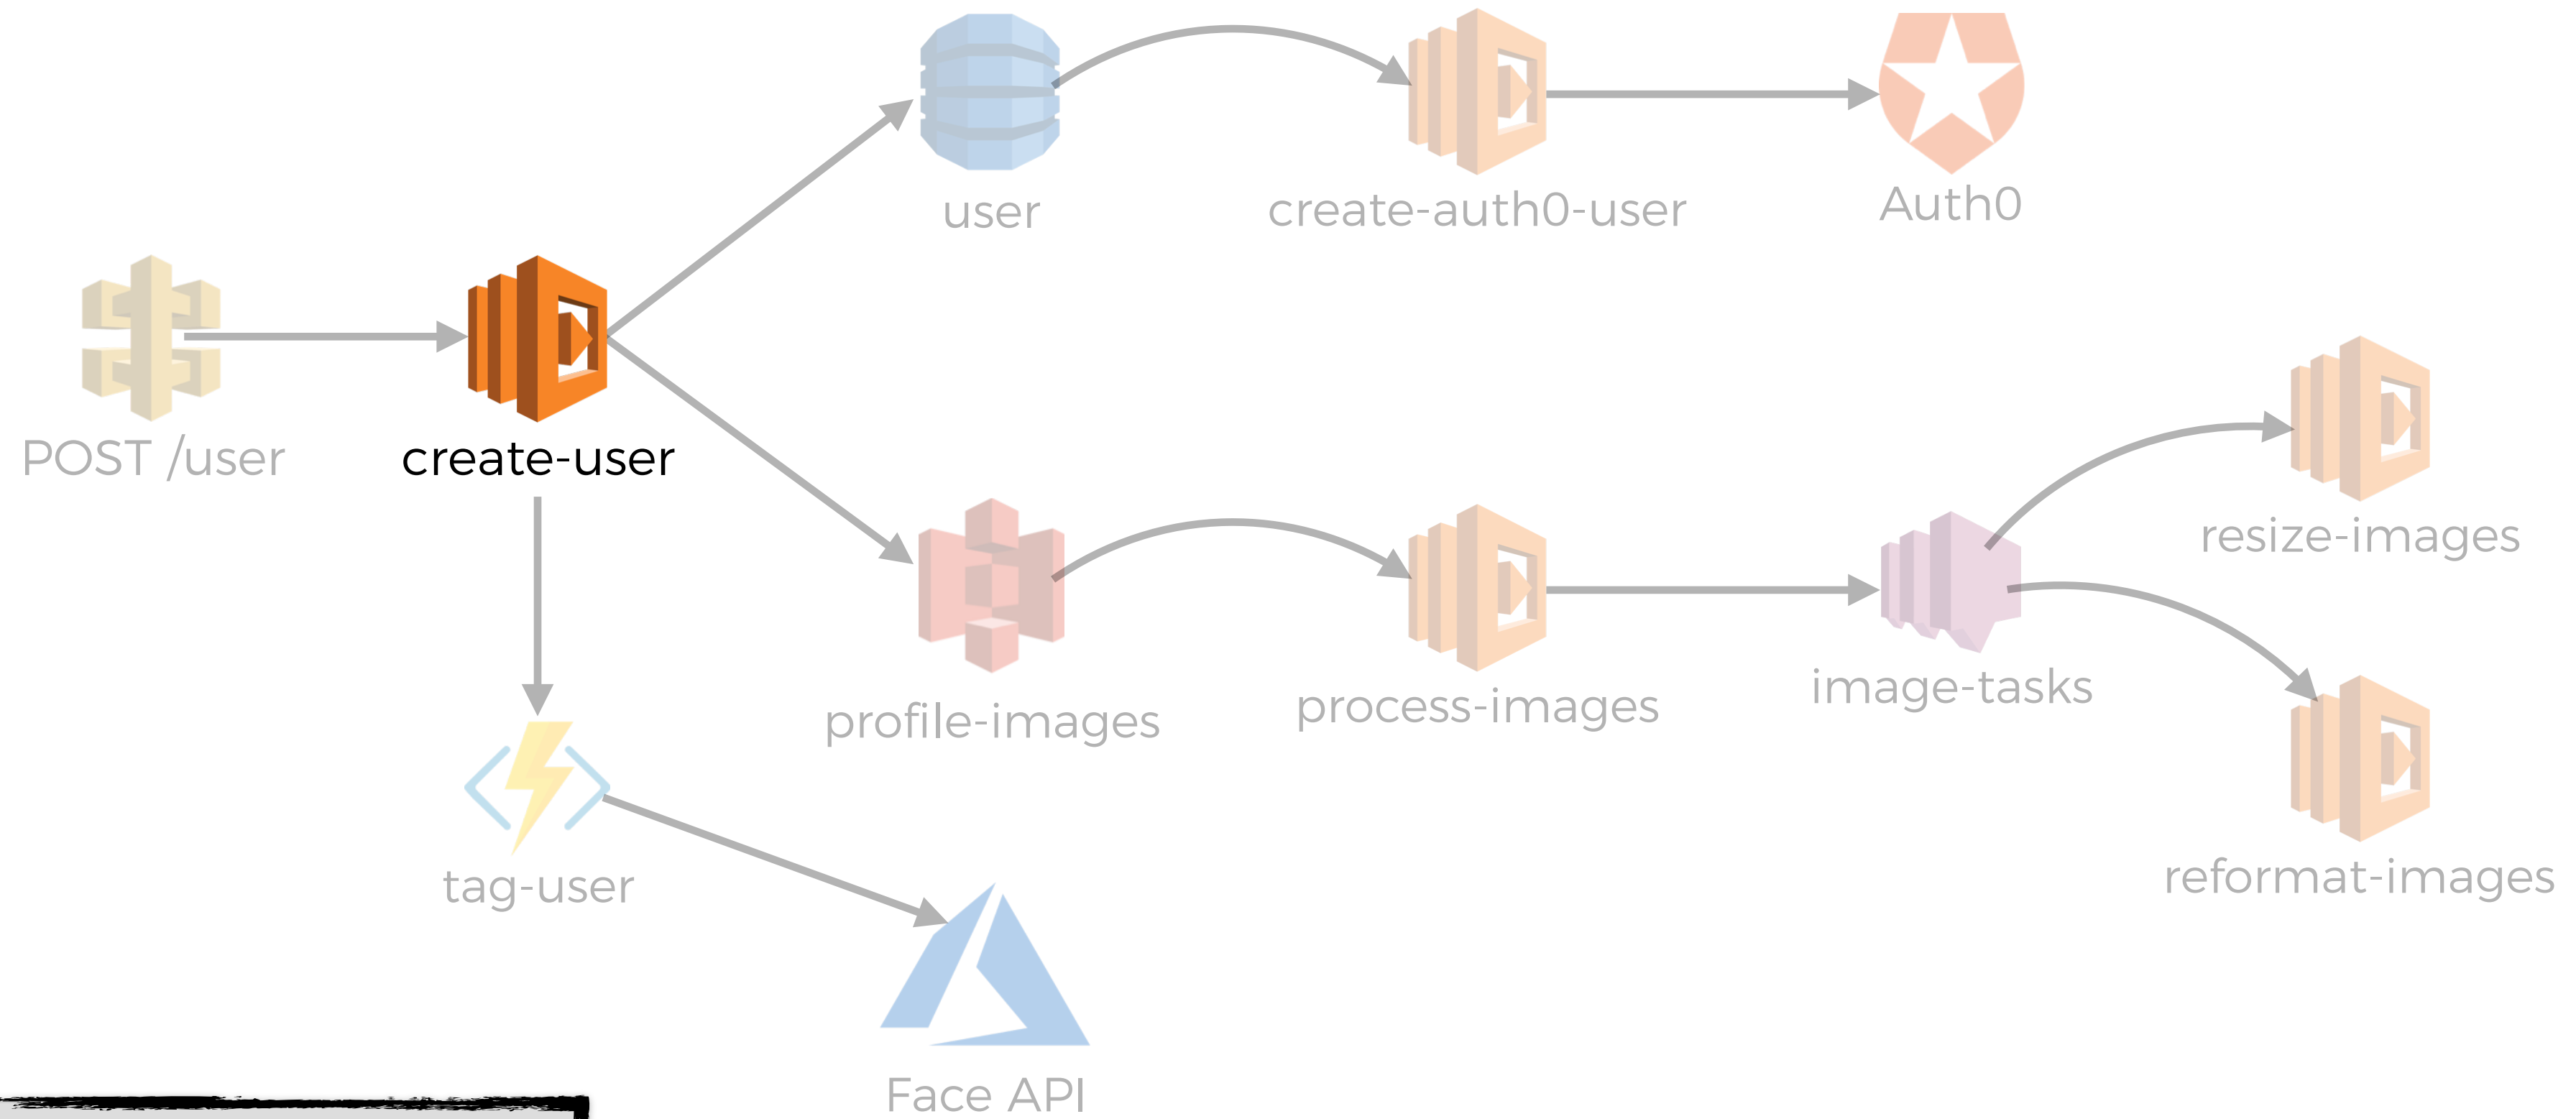
- big-mouth-dev-get-restaurants
 - functions
 - build.sh
 - buildspec.yml
 - package-lock.json
 - package.json
 - README.md
 - seed-restaurants.js
 - template.yml

get-restaurants.js × +

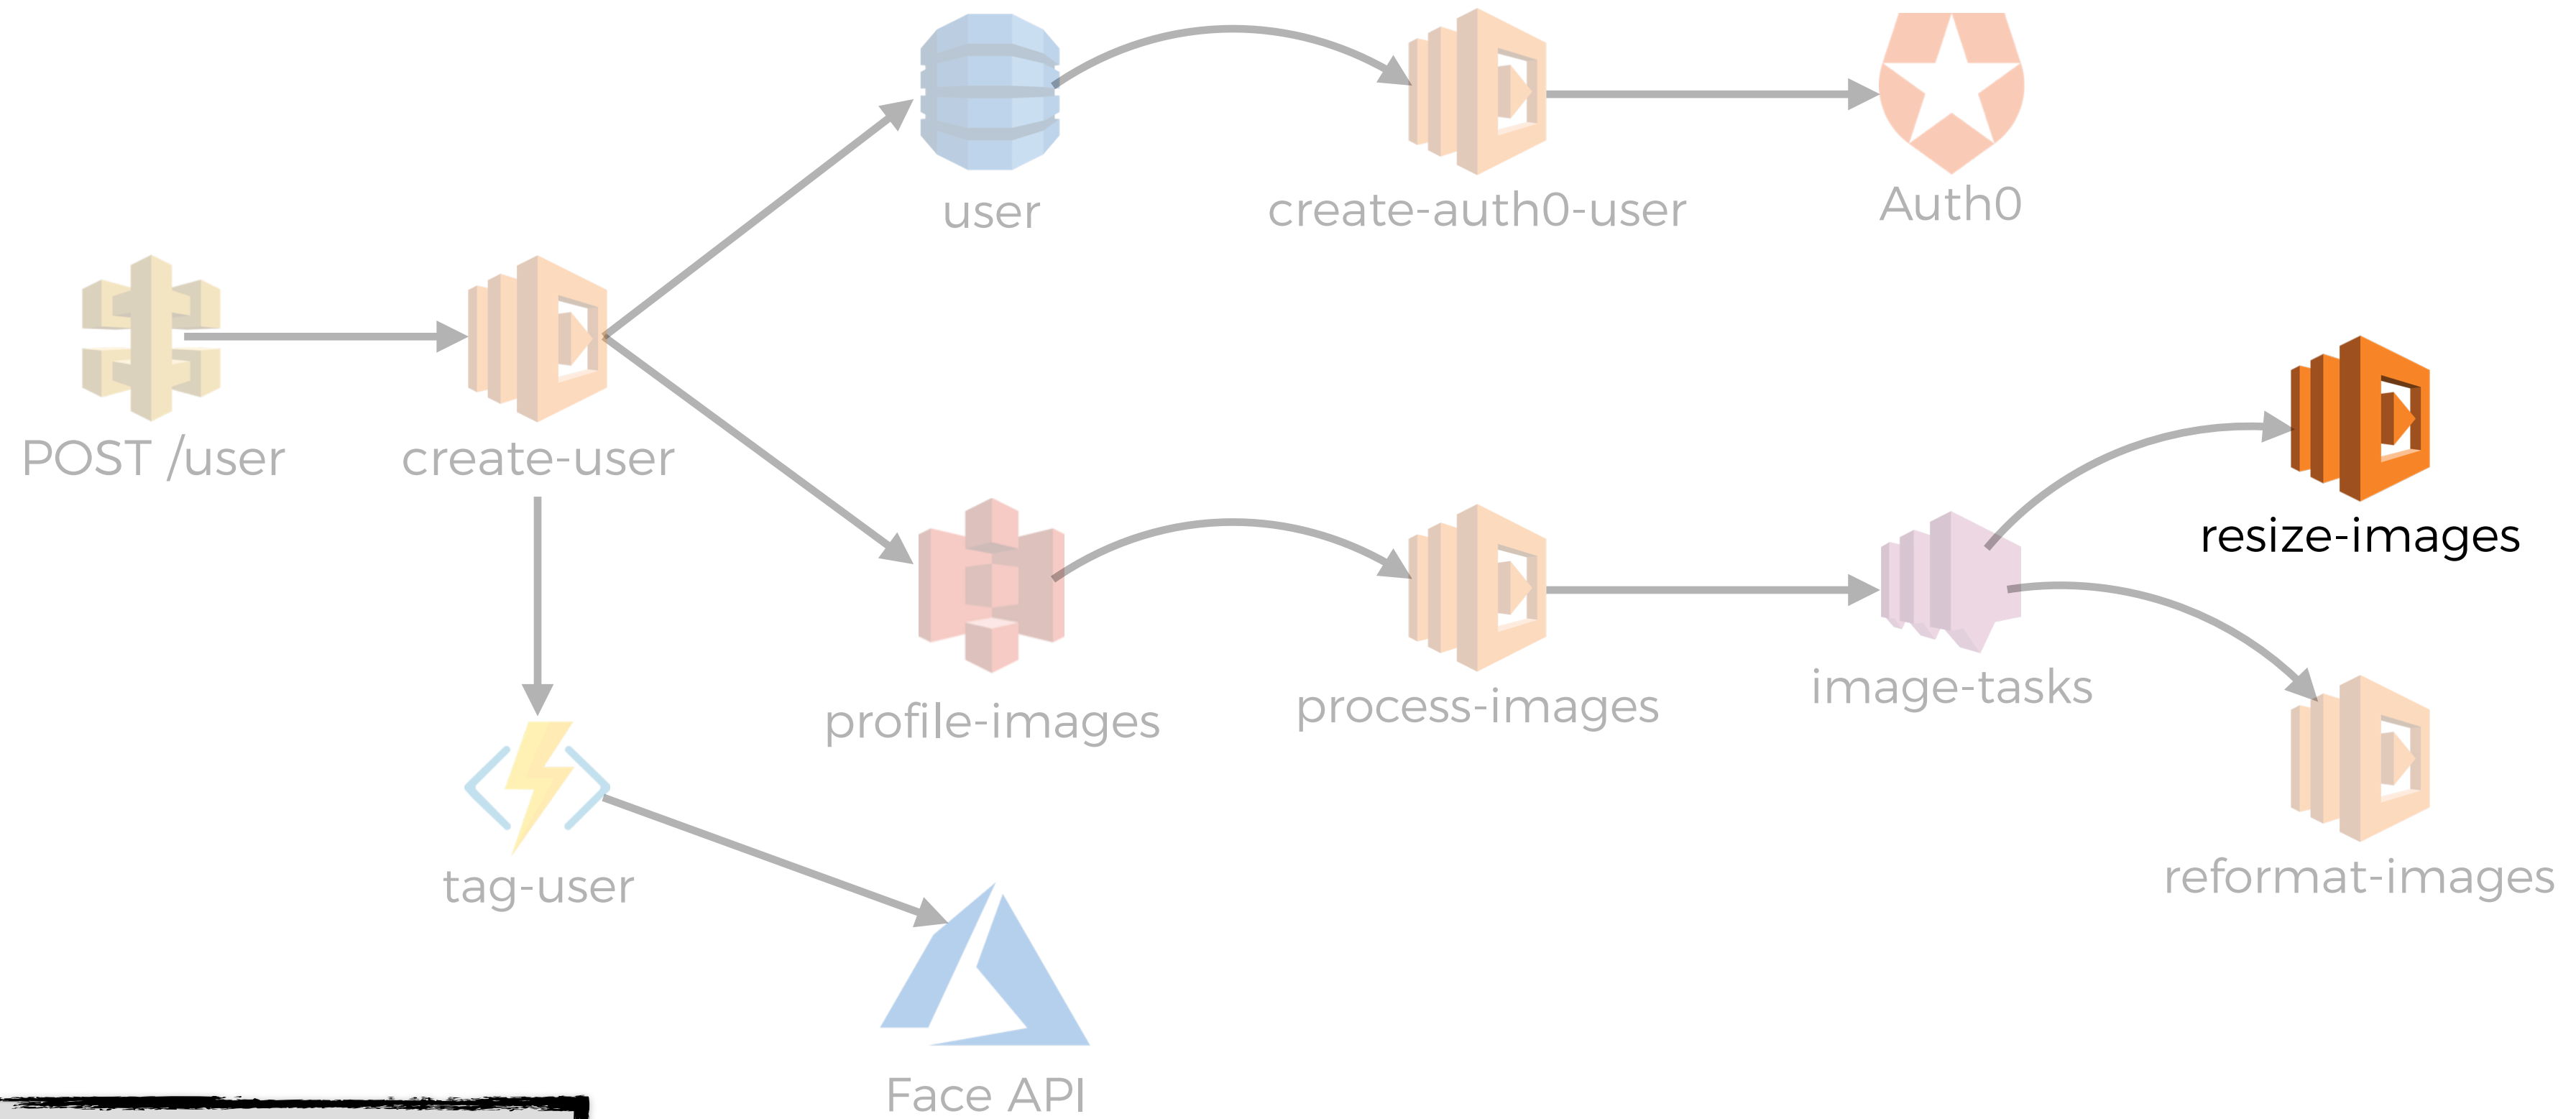
```

1 'use strict';
2
3 const co = require('co');
4 const AWS = require('aws-sdk');
5 const dynamodb = new AWS.DynamoDB.DocumentClient();
6
7 const defaultResults = process.env.defaultResults || 8;
8 const tableName = process.env.restaurants_table;
9
10 function* getRestaurants(count) {
11   let req = {
12     TableName: tableName,
13     Limit: count
14   };
15
16   let resp = yield dynamodb.scan(req).promise();
17   return resp.Items;
18 }
19
20 module.exports.handler = co.wrap(function* (event, context, cb) {
21   let restaurants = yield getRestaurants(defaultResults);
22   let response = {
23     statusCode: 200,
24     body: JSON.stringify(restaurants)
25   };
26
27   cb(null, response);
28 });

```

Logs	Input/Output
	<div style="display: flex; justify-content: space-around;"> input output </div>
	<pre> { "body": "{ \"username\": \"theburningmonk\"}", "resource": "/user", "requestContext": { "resourceId": "123456", "apiId": "1234567890", "resourcePath": "/user", } } </pre>



Logs	Input/Output
------	--------------

input	output
-------	--------

```

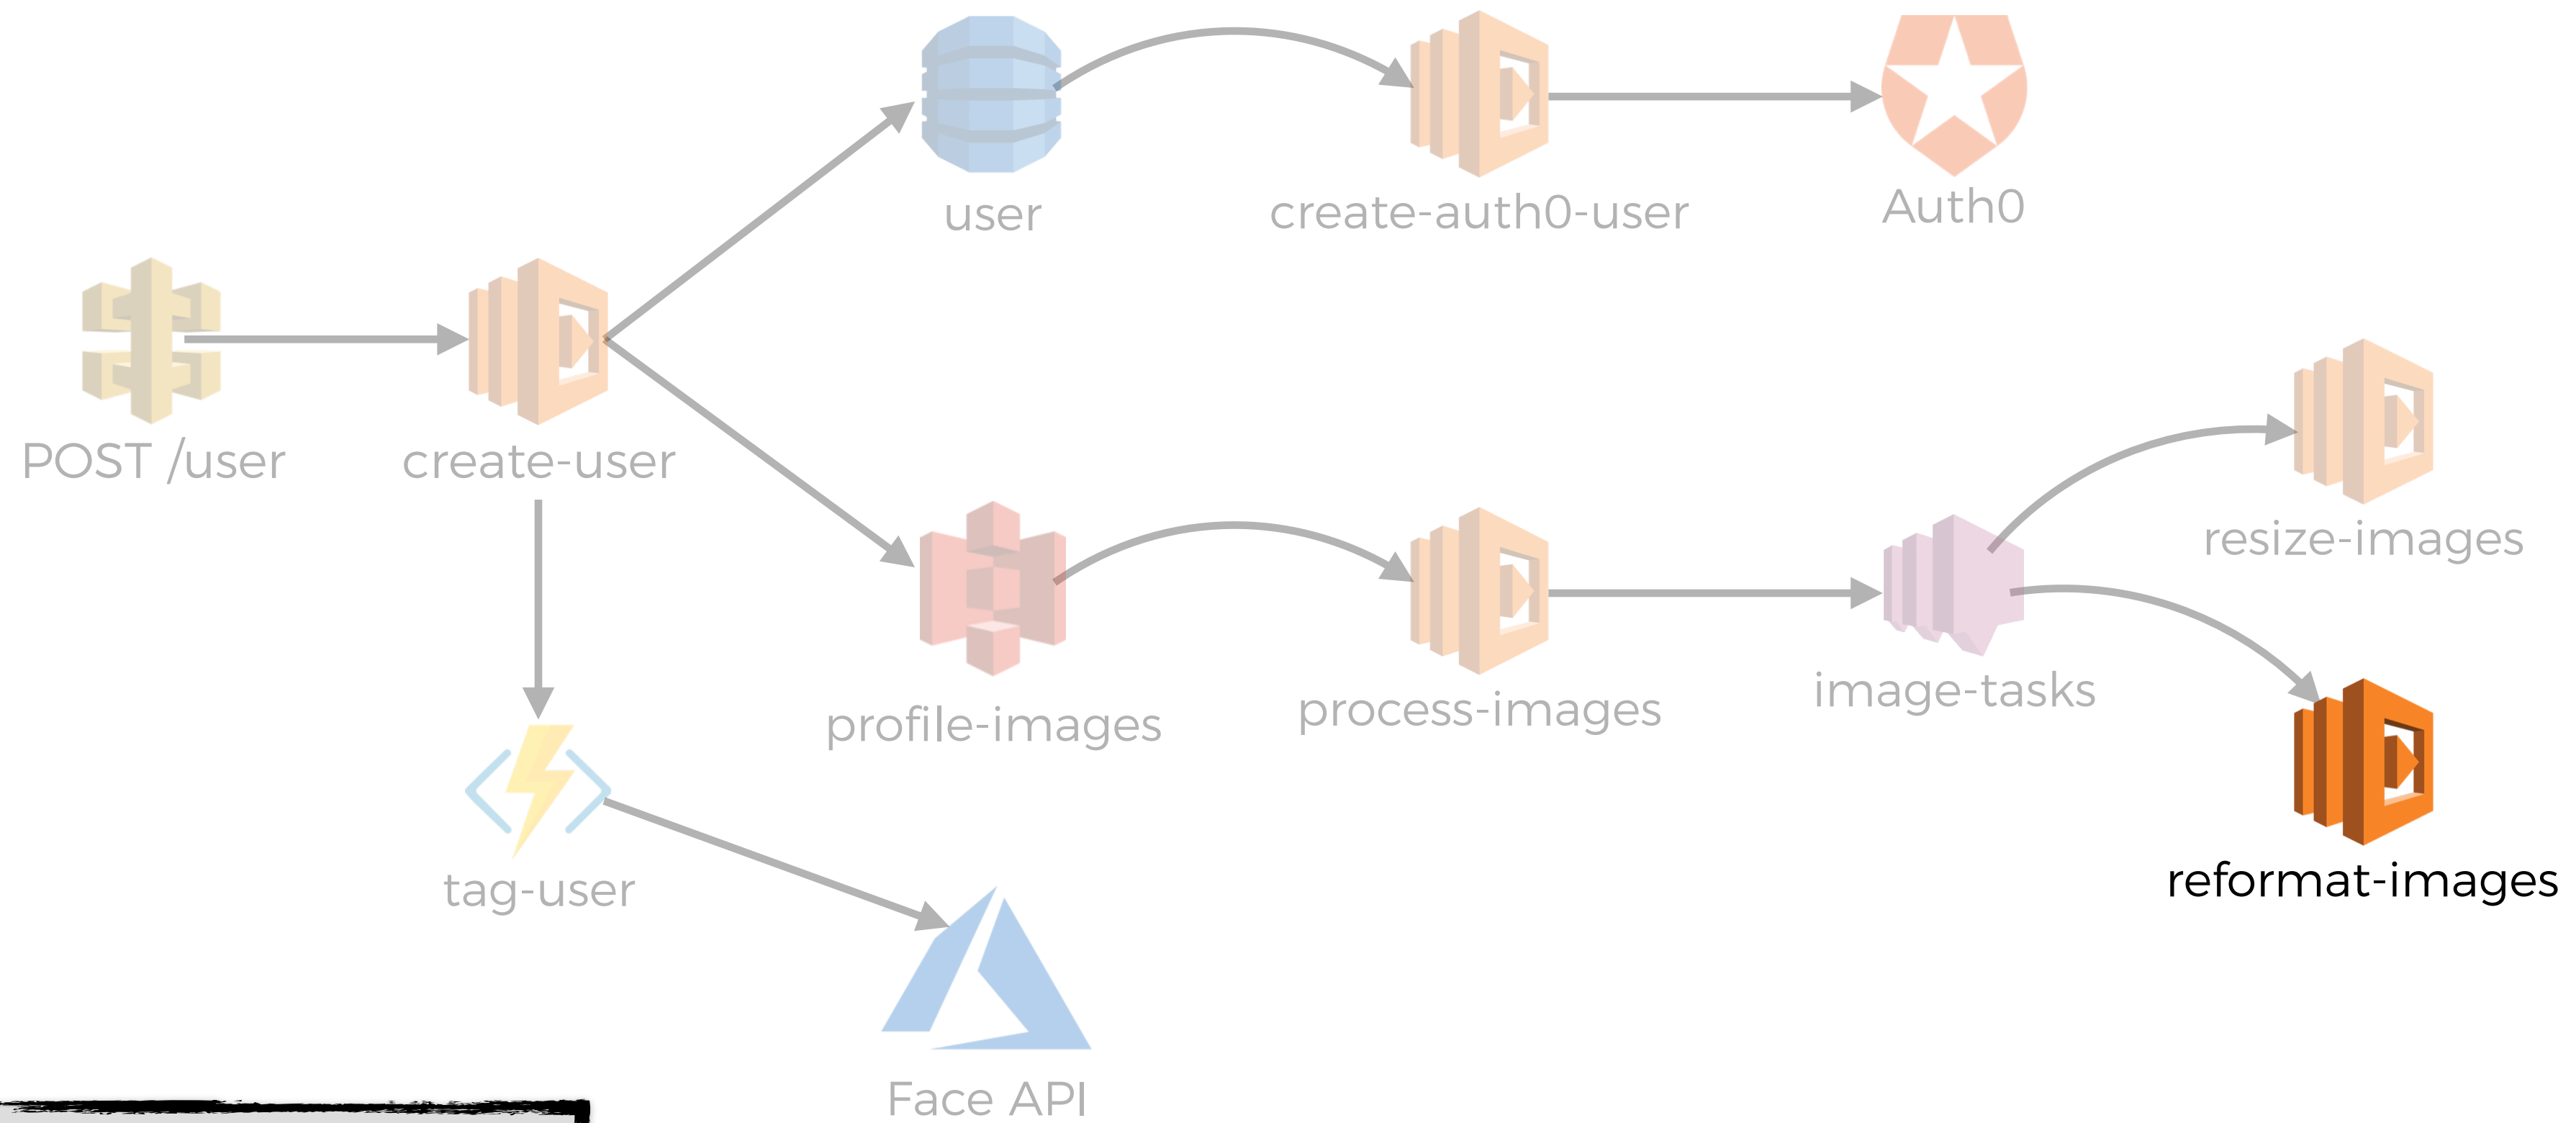
{ "Records": [
  { "Sns": {
    "Type": "Notification",
    "MessageId": "...",
    "TopicArn": "...",
    "Message": "...",
    "Timestamp": "2018/01/25 20:51:24.215",
  }
}

```

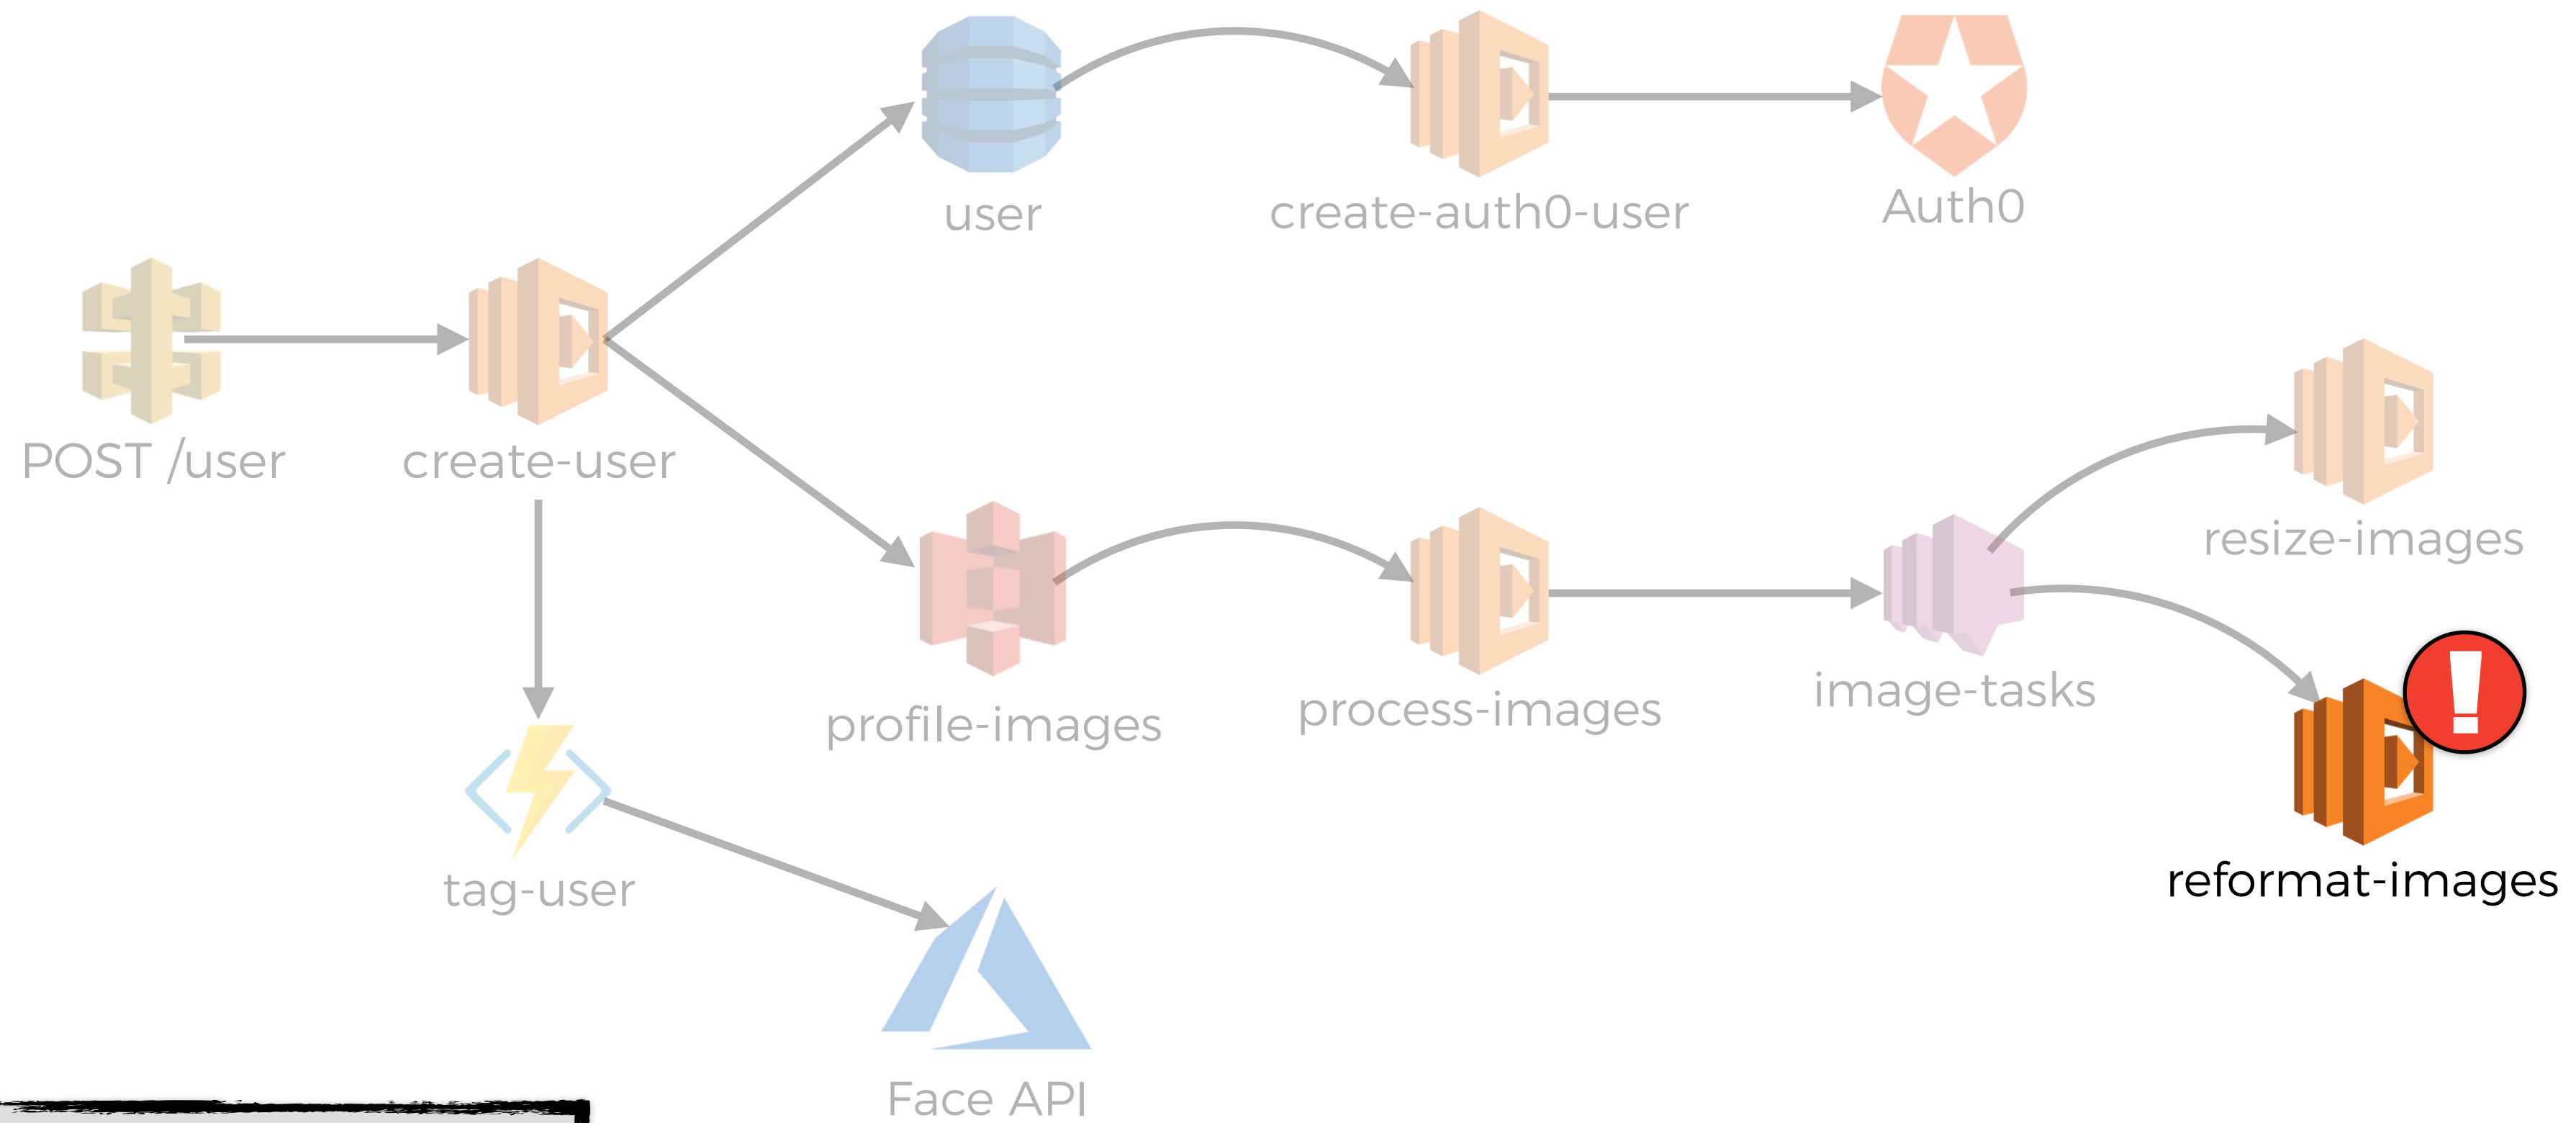
```

{
  "error": null,
  "result": "OK"
}

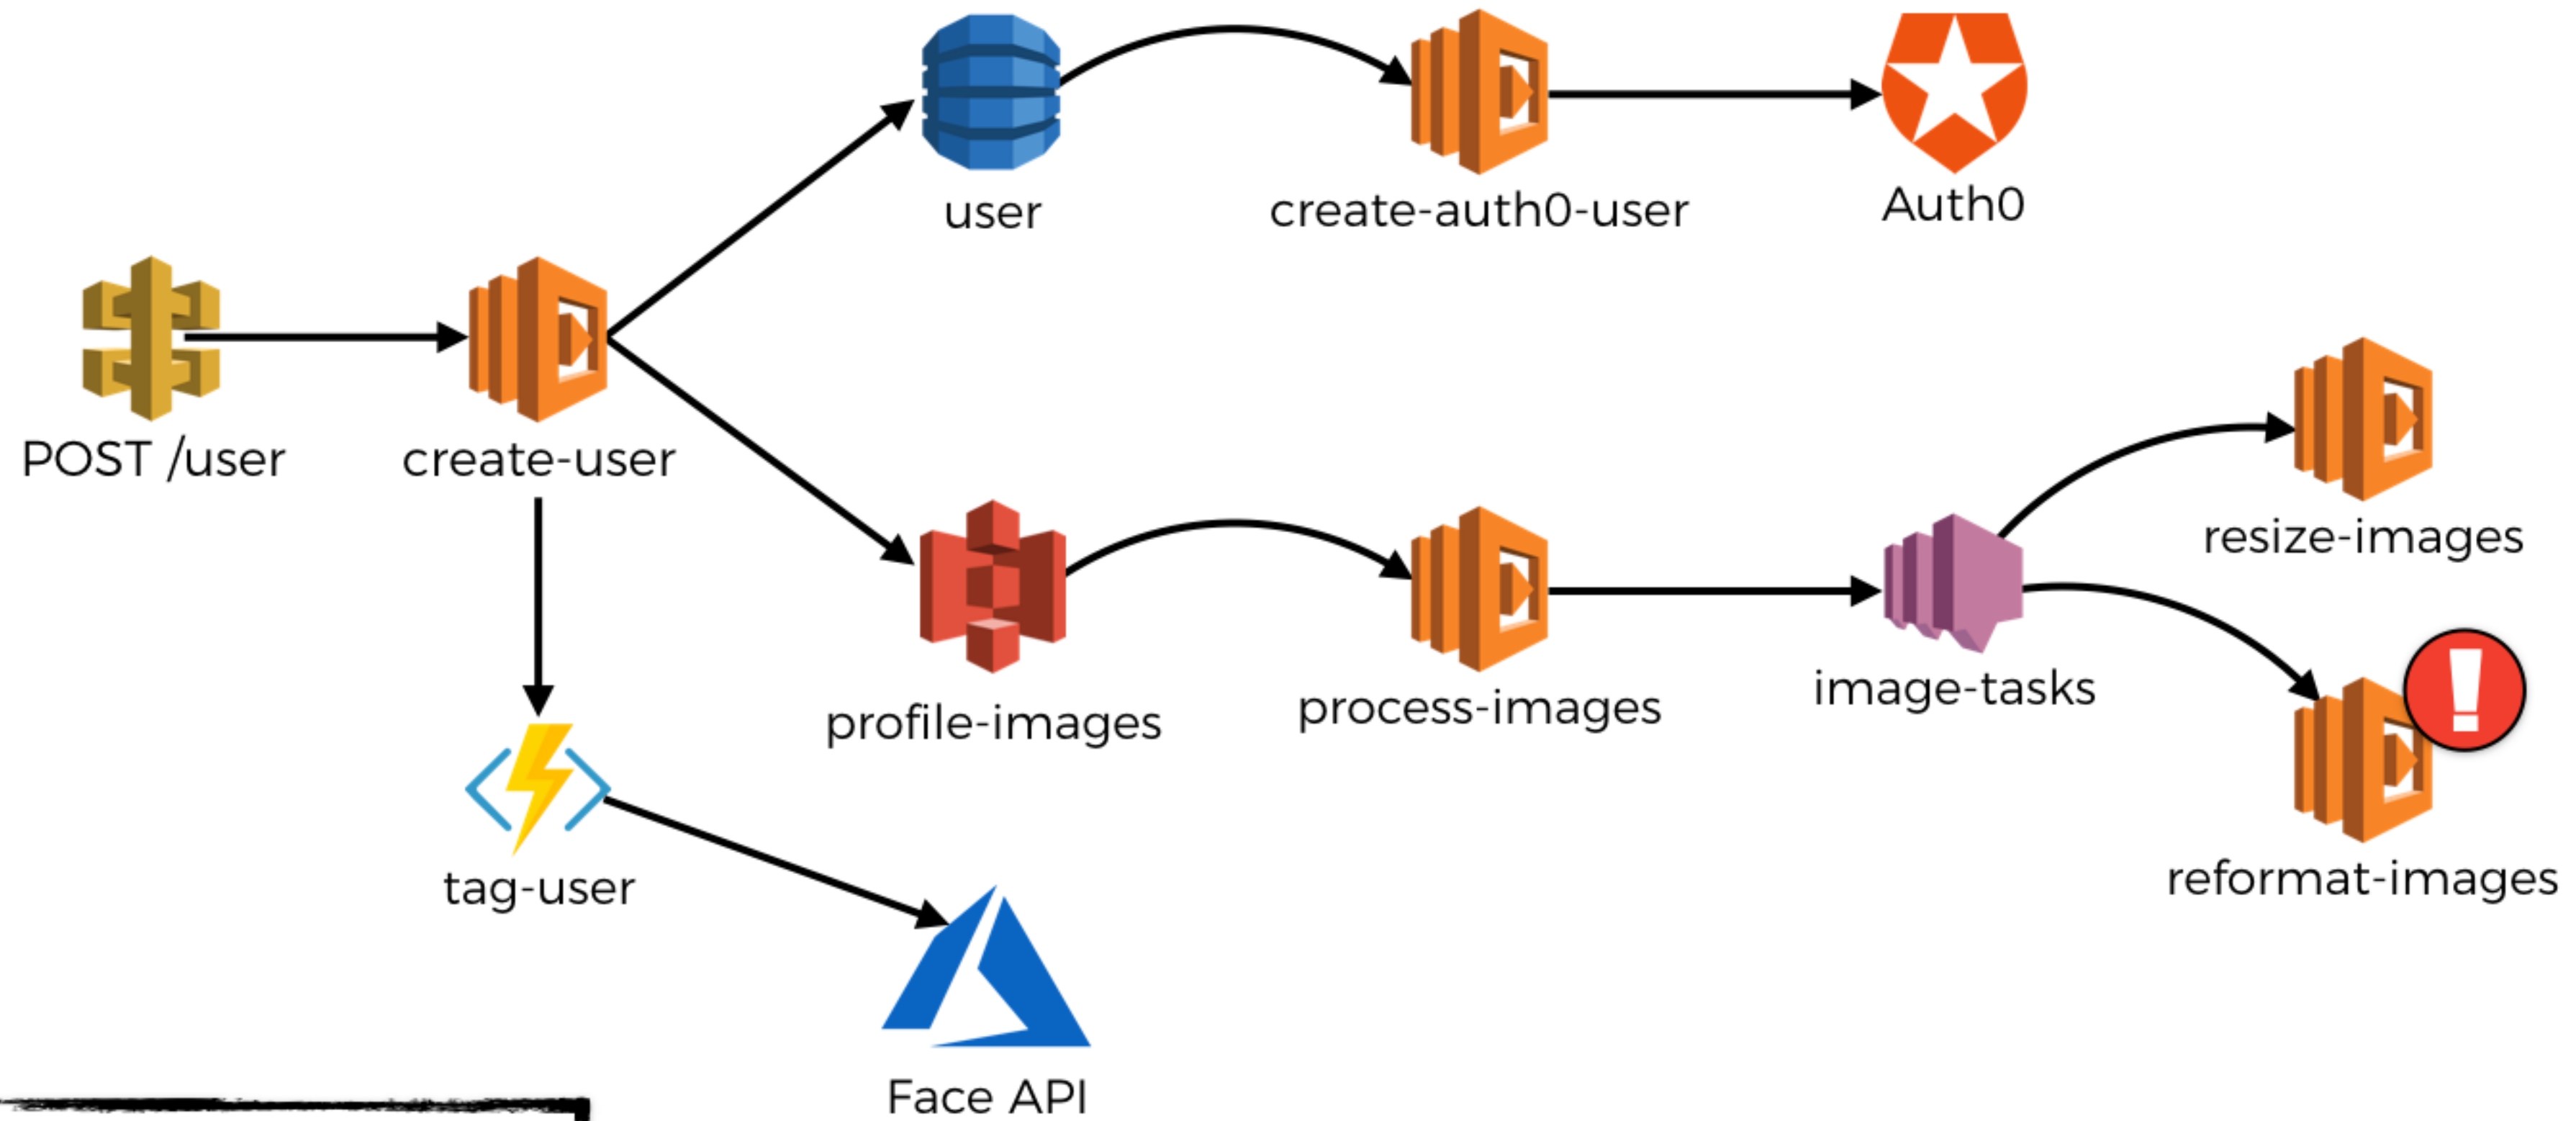
```

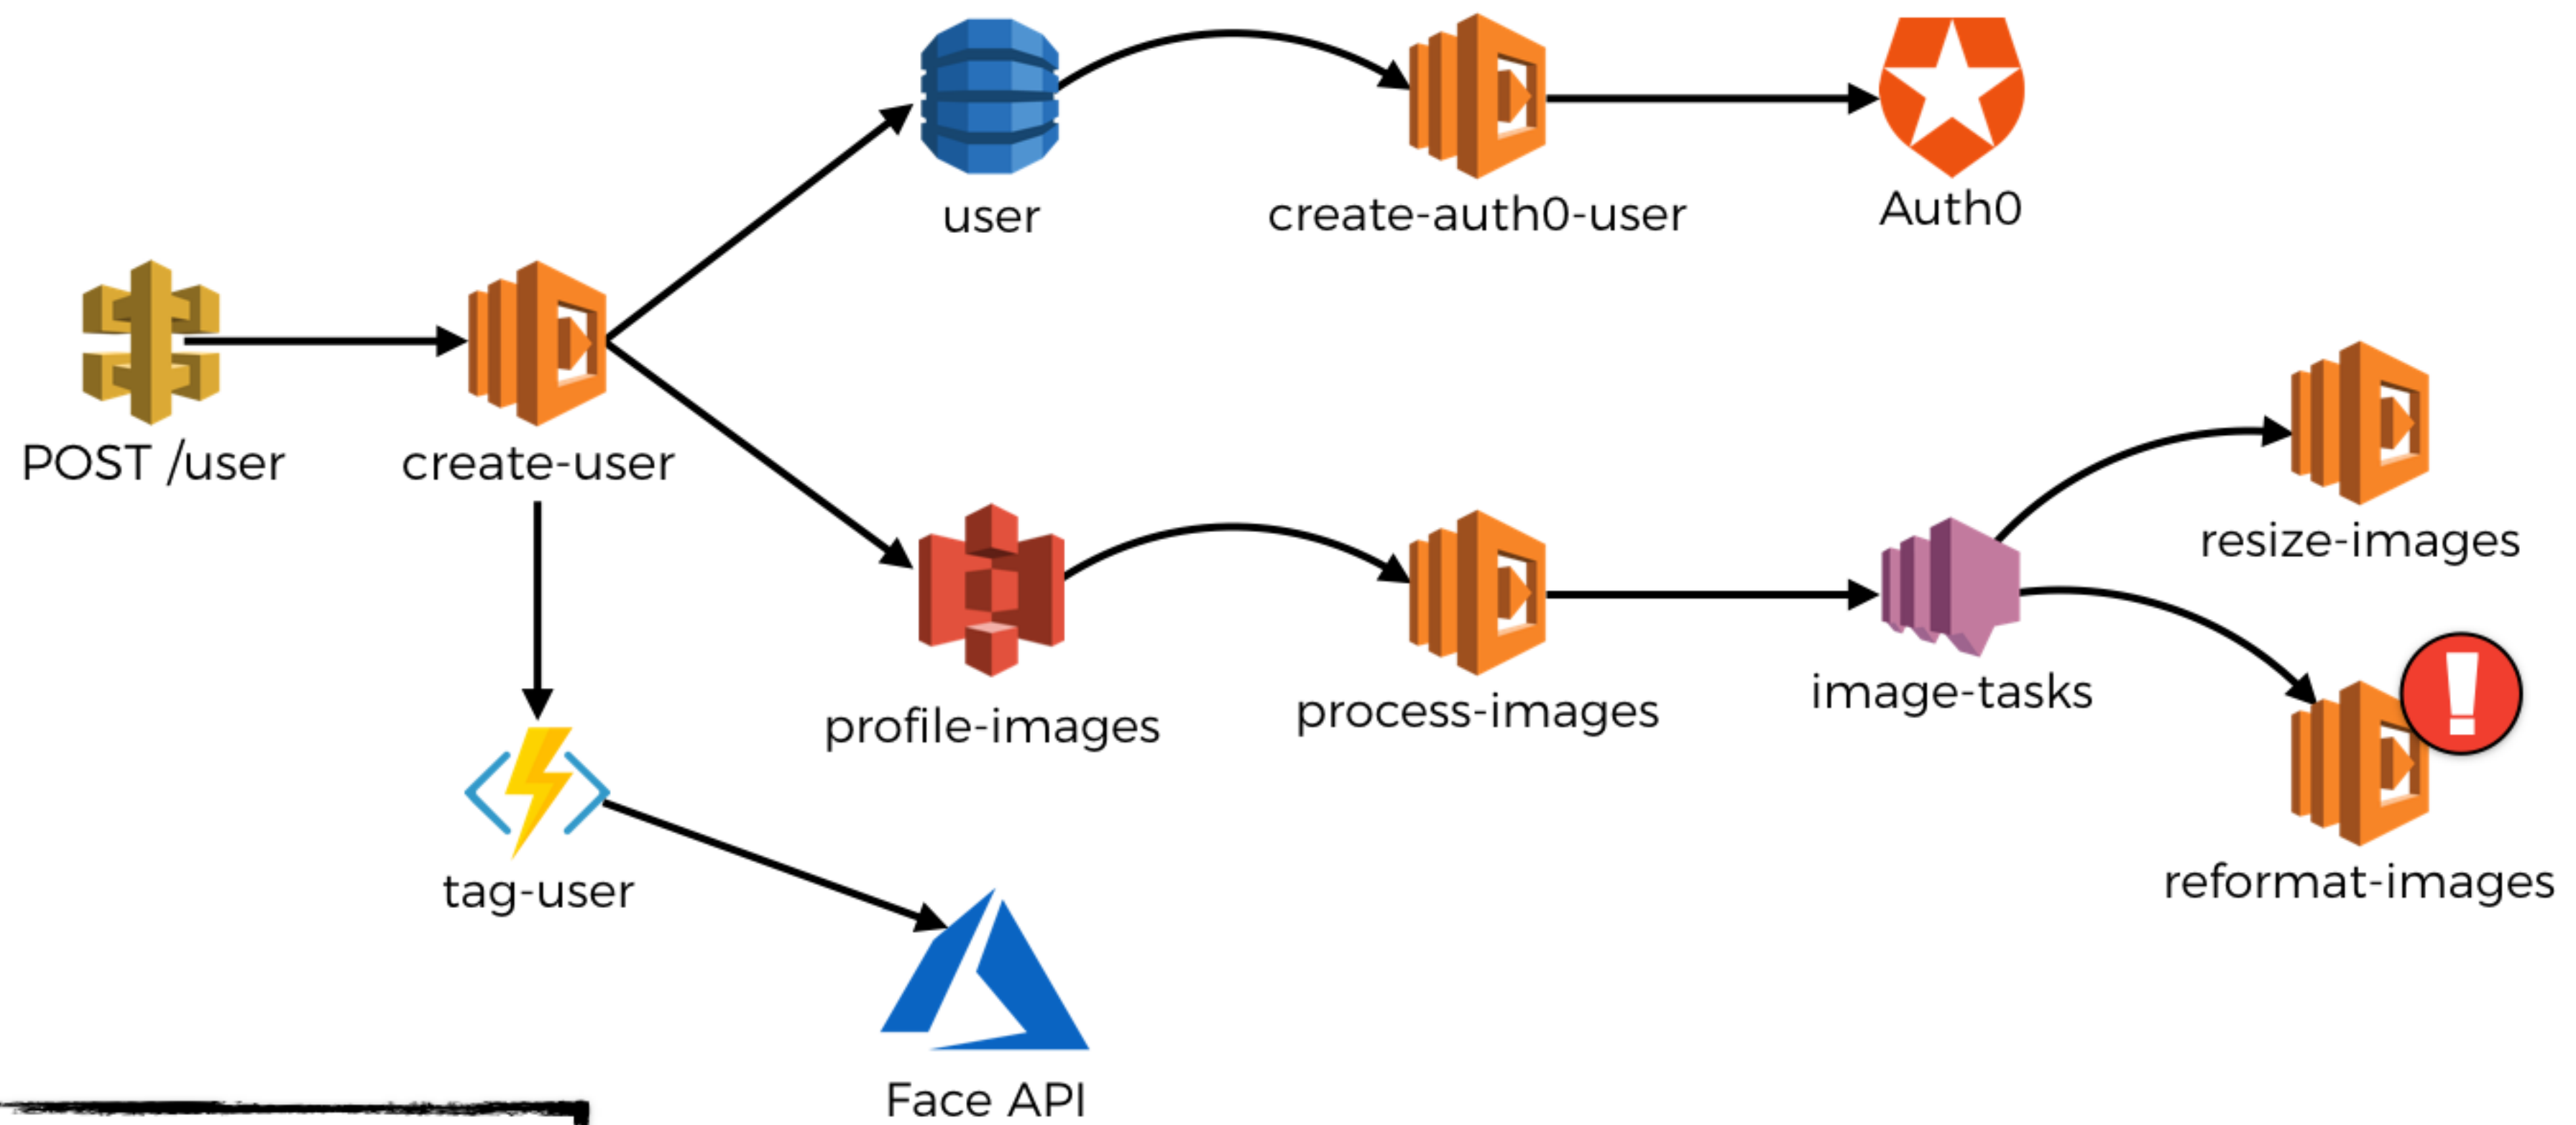
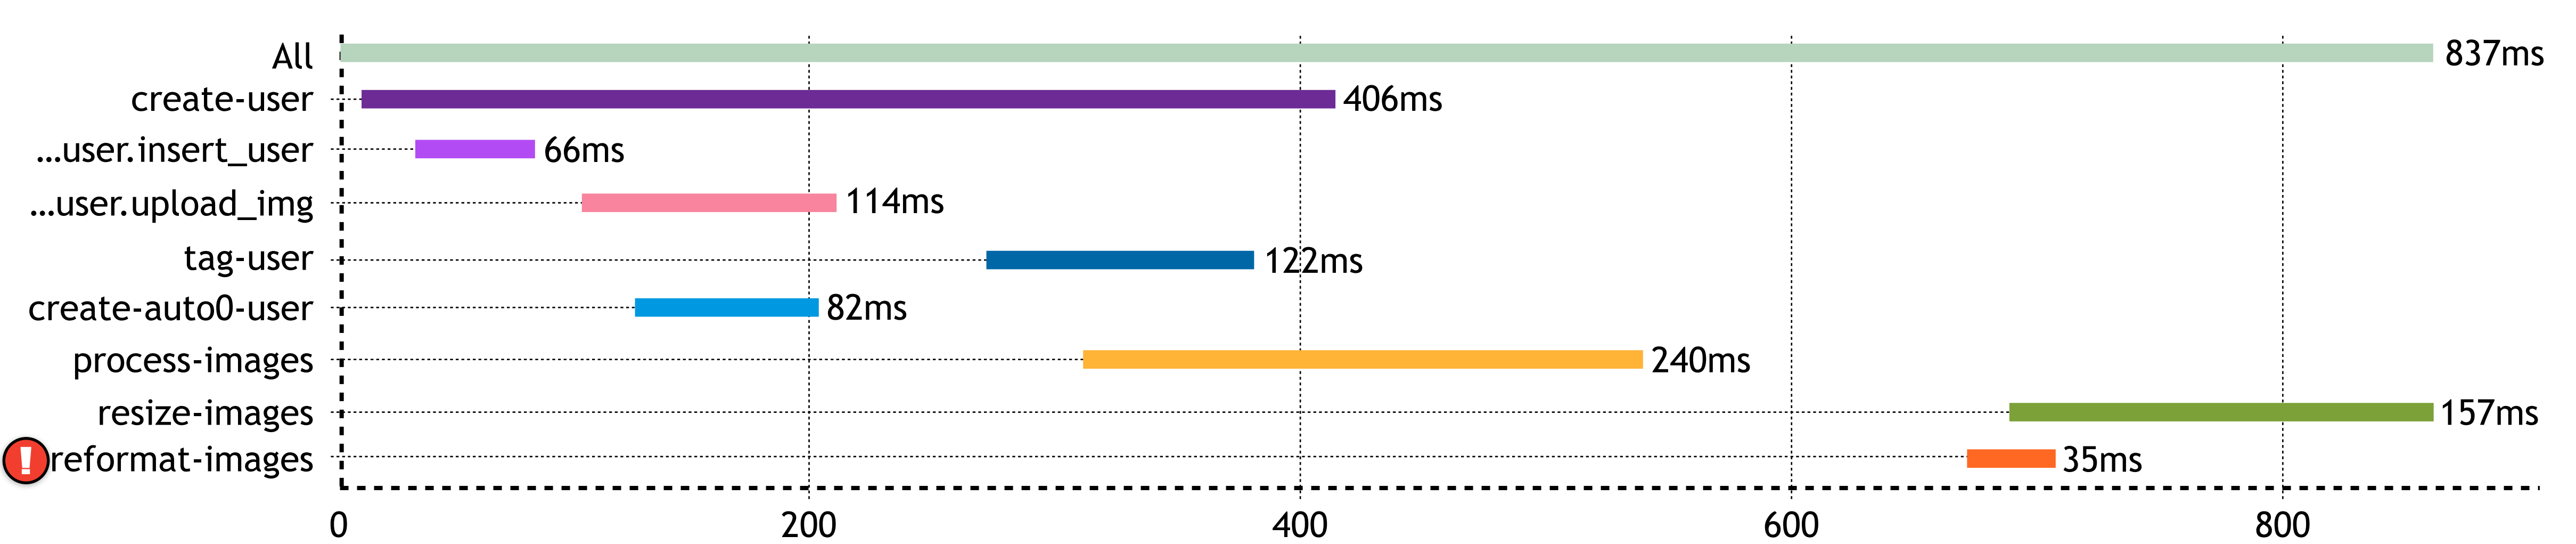
Logs	Input/Output
	input
<pre> { "Records": [{ "Sns": { "Type": "Notification", "MessageId": "...", "TopicArn": "...", "Message": "...", "Timestamp": "2018/01/25 20:51:24.215", </pre>	error
	<pre> [com.spaceape.dragon.handler.ReformatProfileImageHandle r] Null reference exception *java.lang.NullReferenceException: ... * at ... * at ... * at ... </pre>



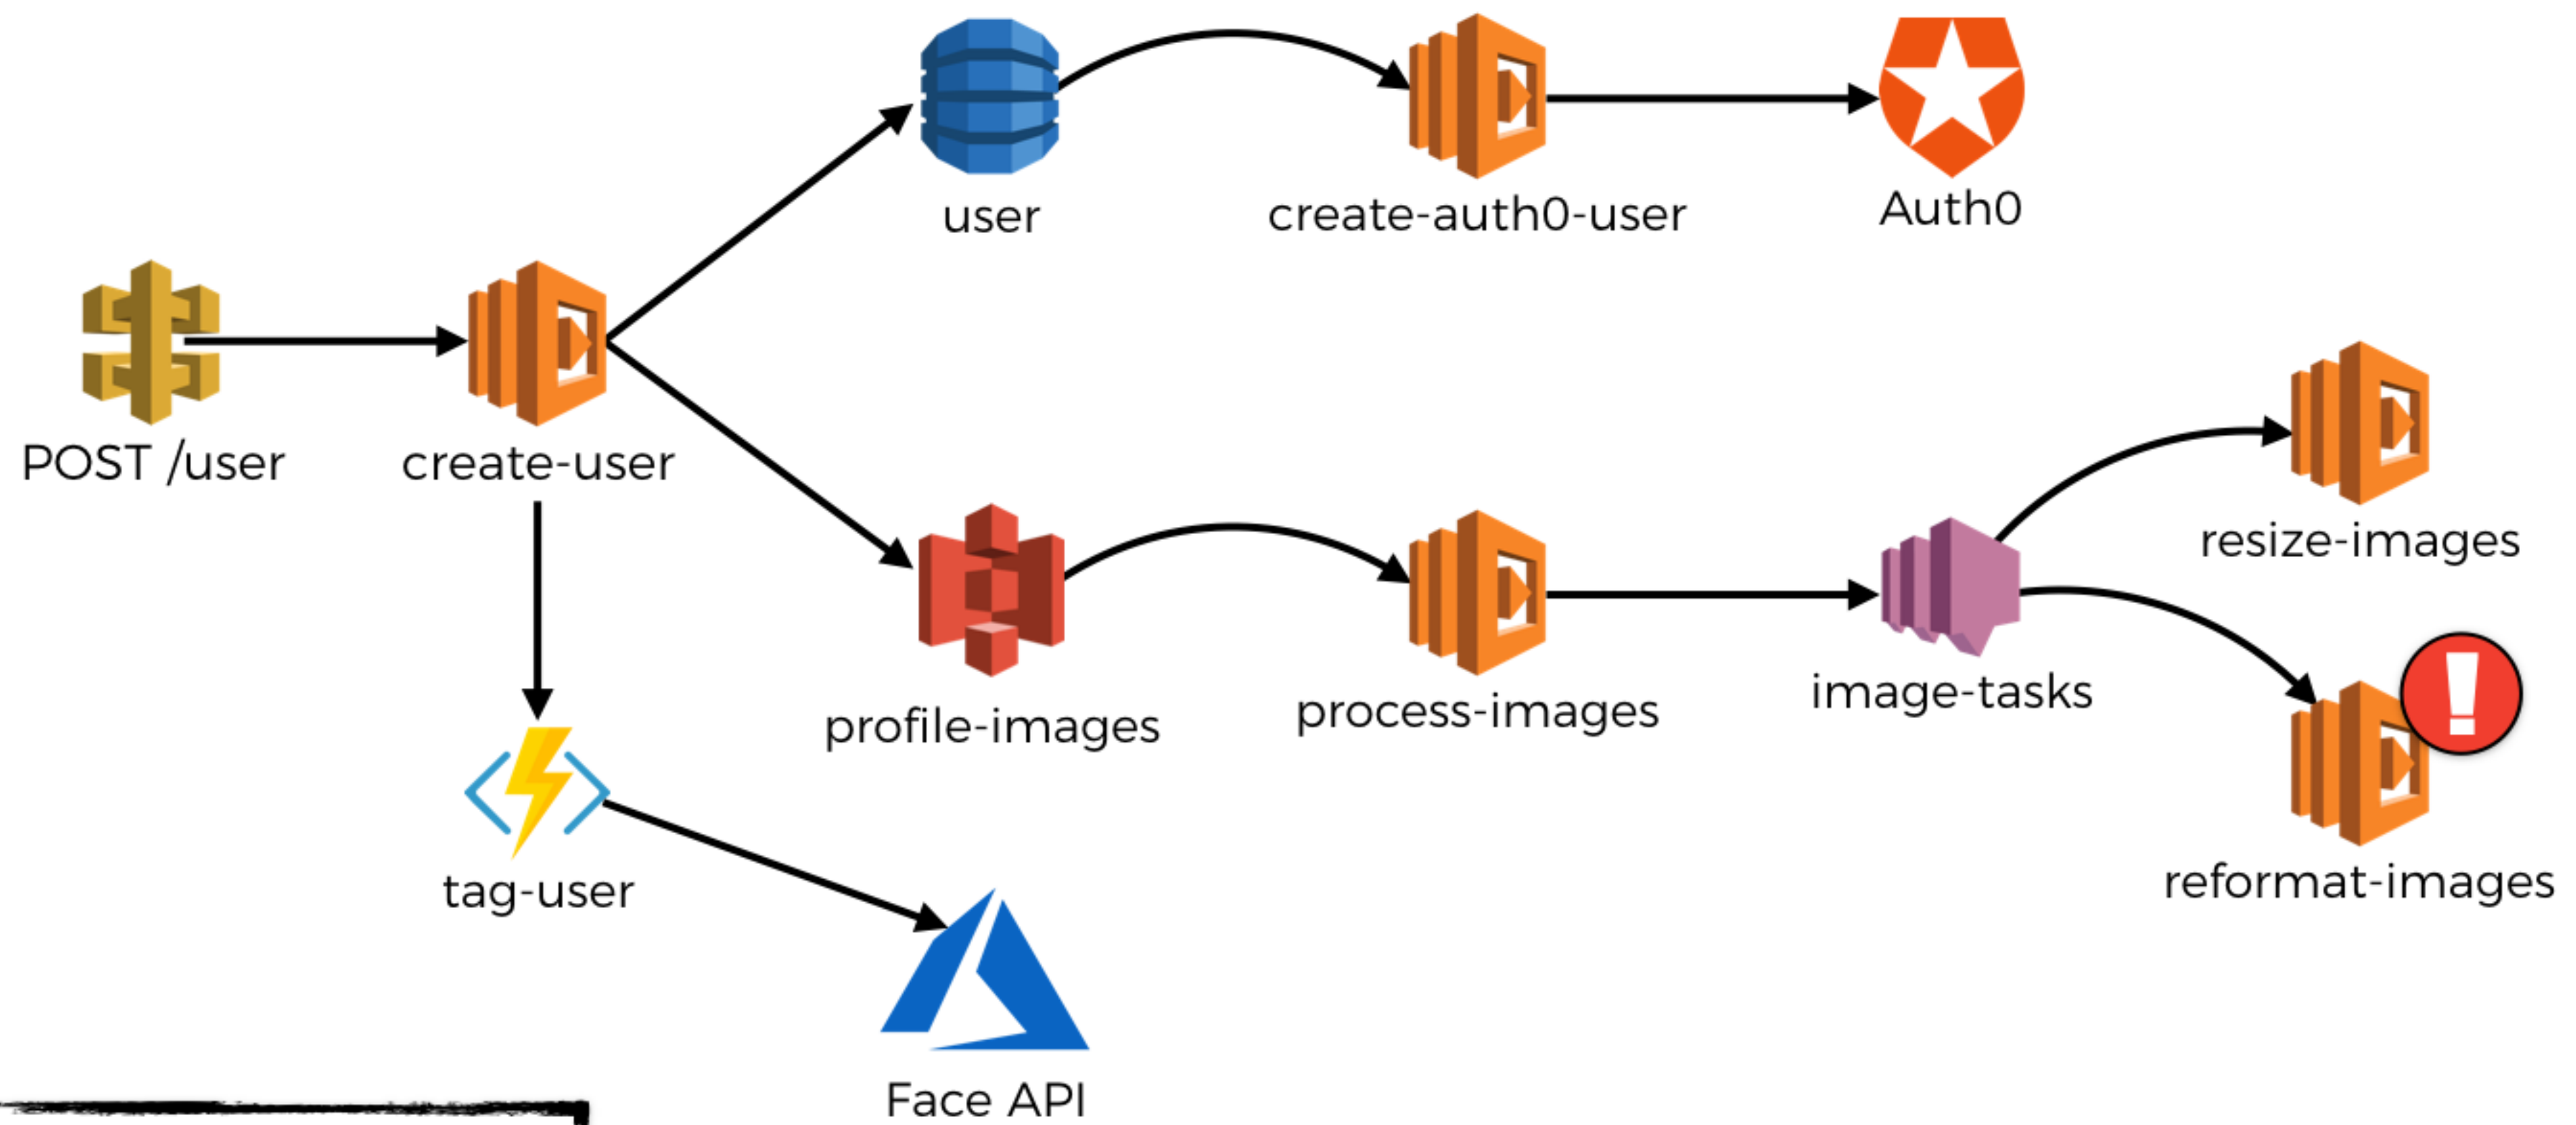
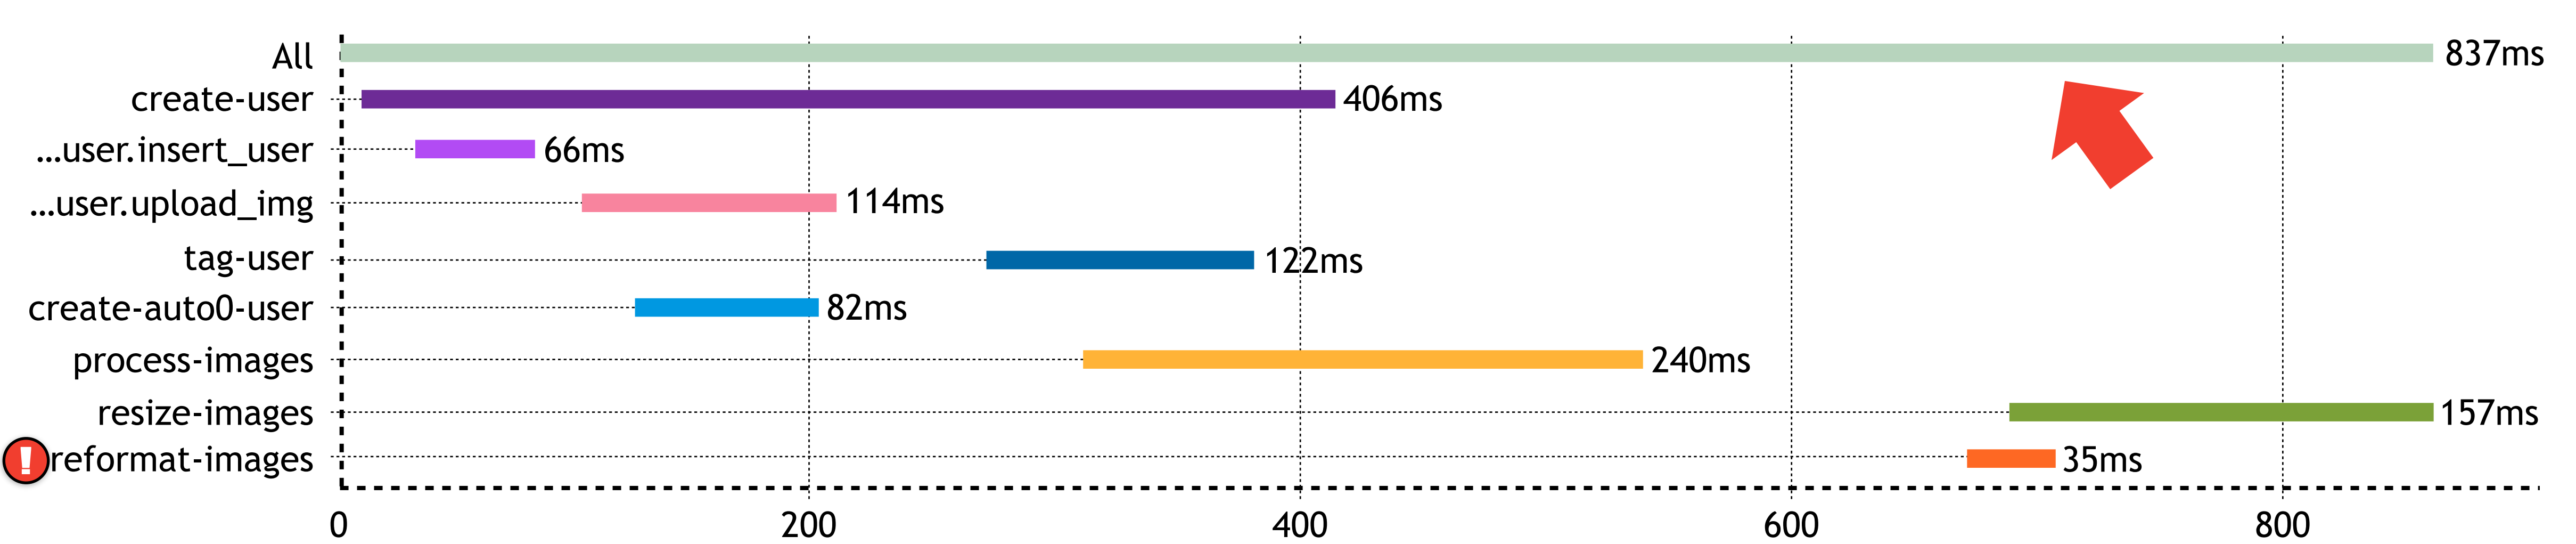
Logs	Input/Output
	input
<pre> { "Records": [{ "Sns": { "Type": "Notification", "MessageId": "...", "TopicArn": "...", "Message": "...", "Timestamp": "2018/01/25 20:51:24.215", </pre>	error
	<pre> [com.spaceape.dragon.handler.ReformatProfileImageHandle r] Null reference exception *java.lang.NullReferenceException: ... * at ... * at ... * at ... </pre>



Logs		Input/Output		
timestamp	component	level	message	
2018/01/25 20:51:23.188	POST /user	debug	incoming request...	
2018/01/25 20:51:23.201	create-user	debug	saving user [theburningmonk] in the [user] table...	
2018/01/25 20:51:23.215	create-user	debug	saved user [theburningmonk] in the [user] table	
2018/01/25 20:51:23.521	tag-user	debug	tagging user [theburningmonk] with Azure Face API...	

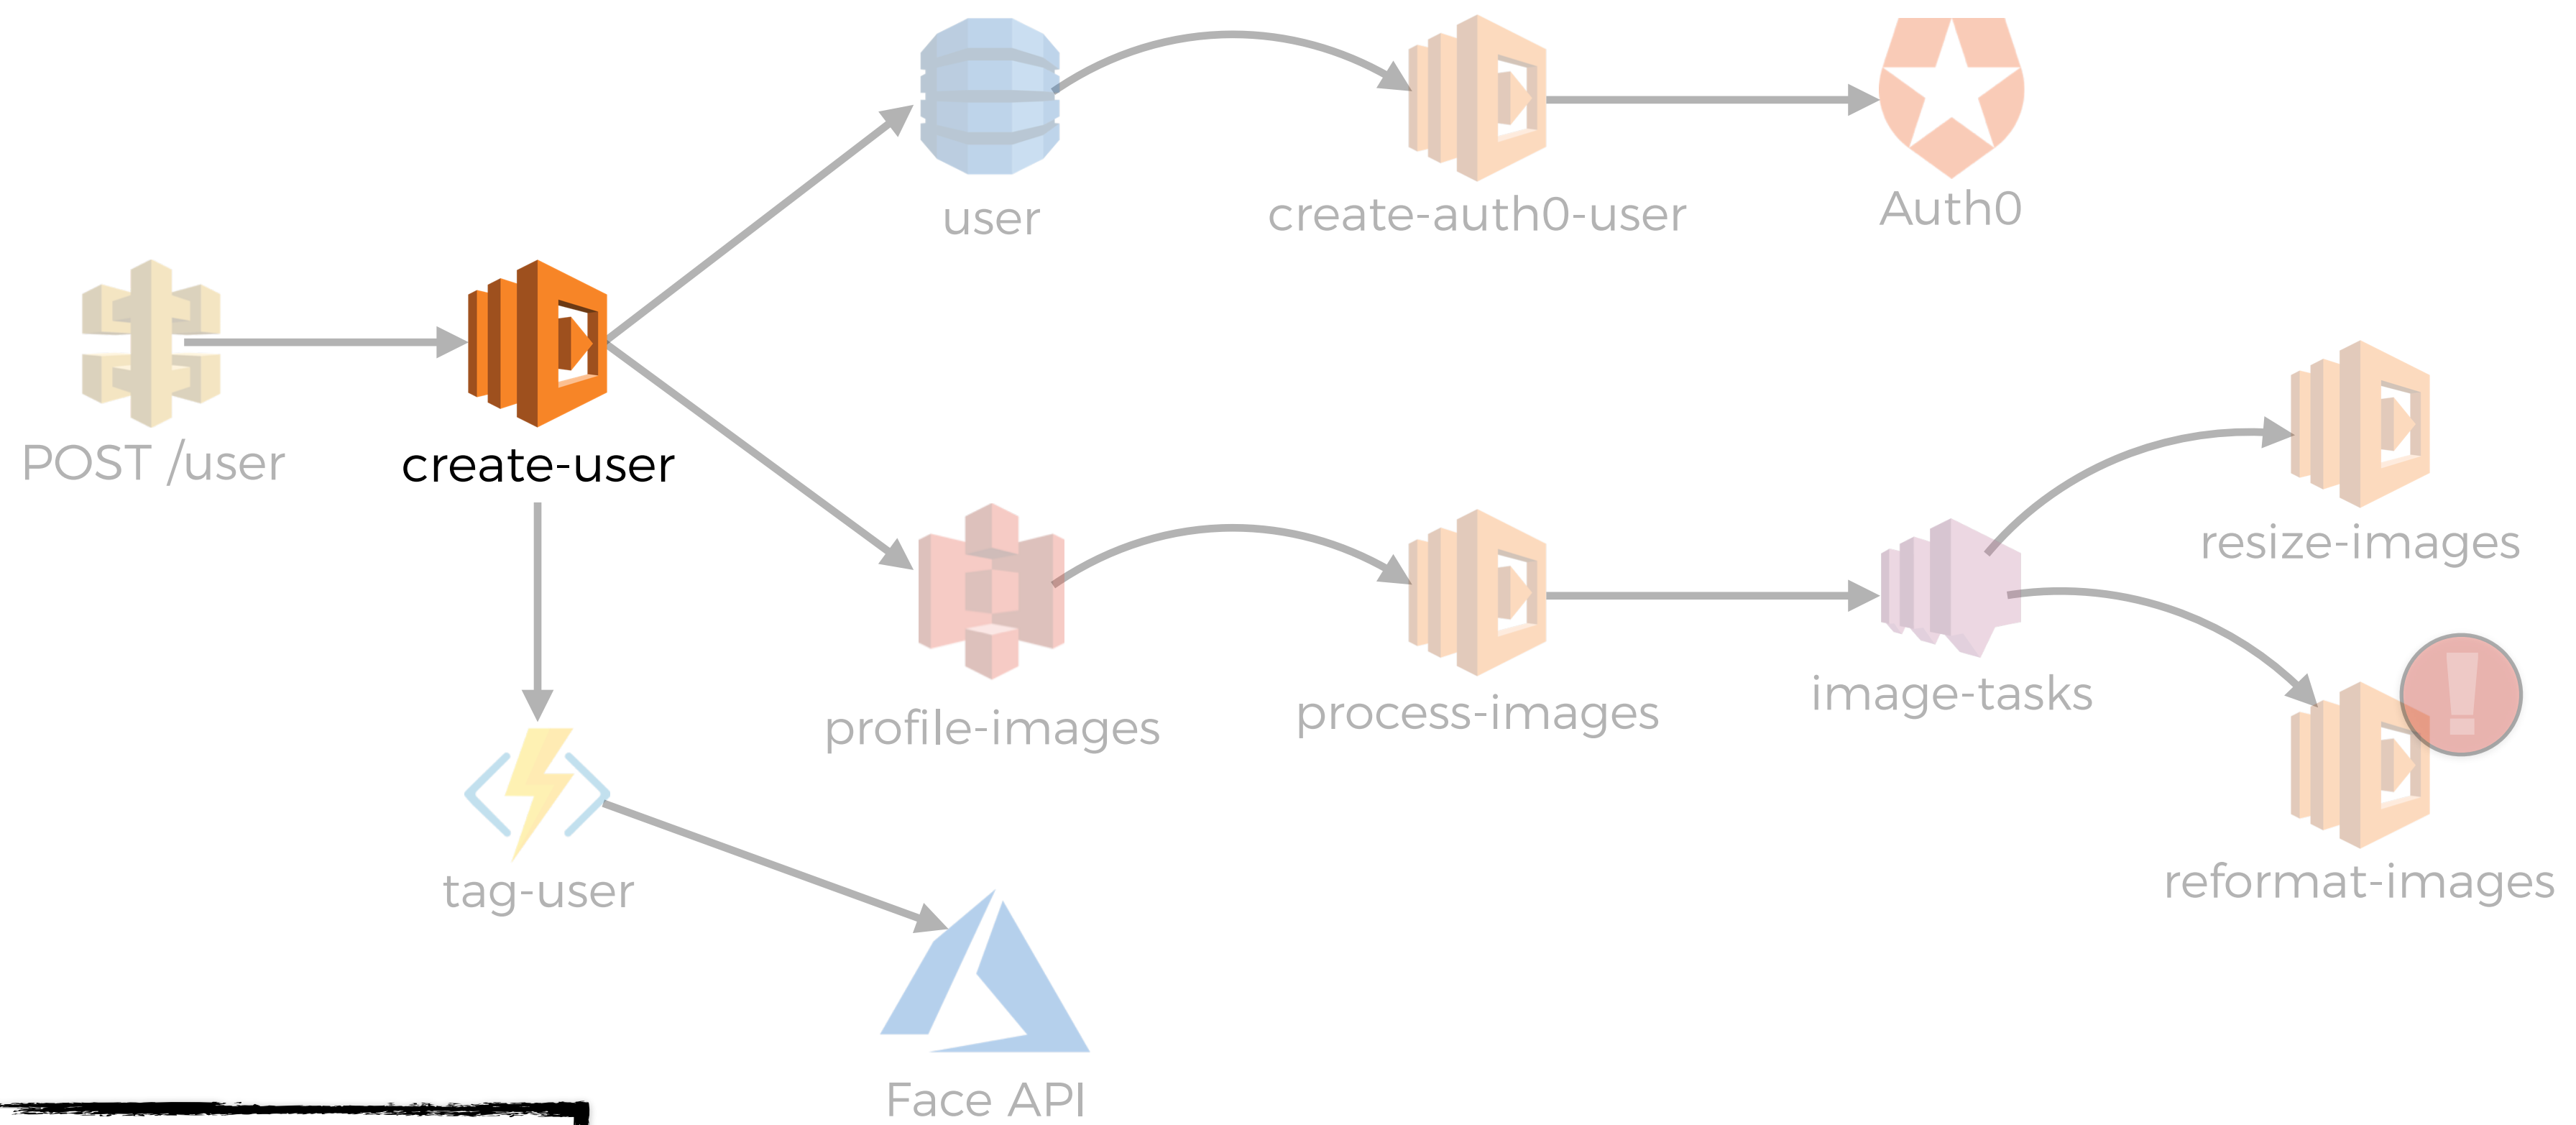
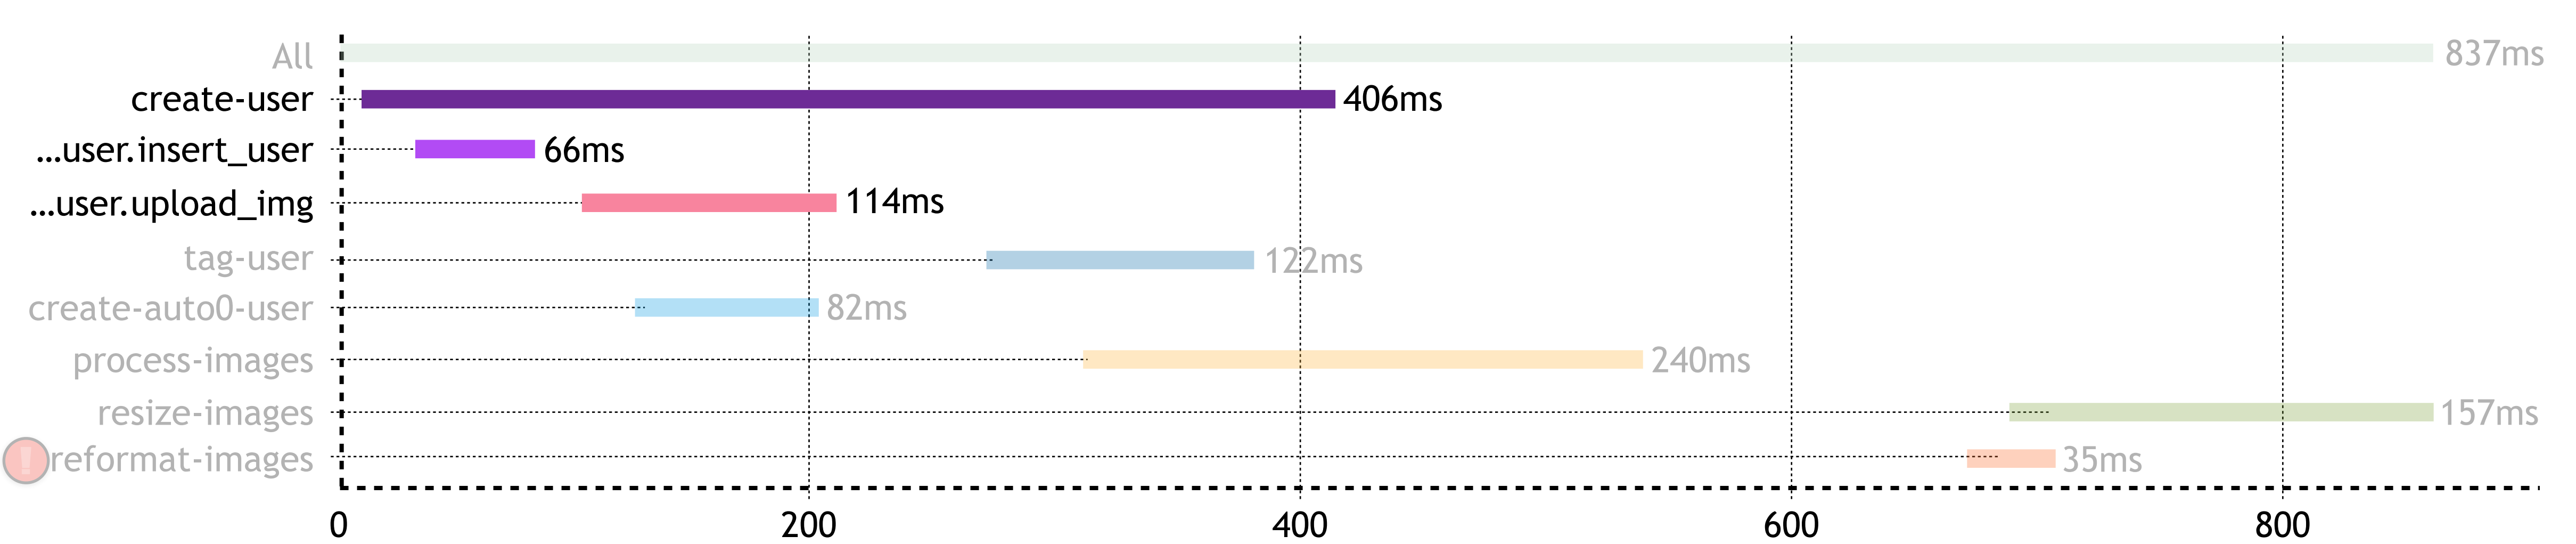


Logs	Input/Output
timestamp	component
	level
	message



Logs	Input/Output
------	--------------

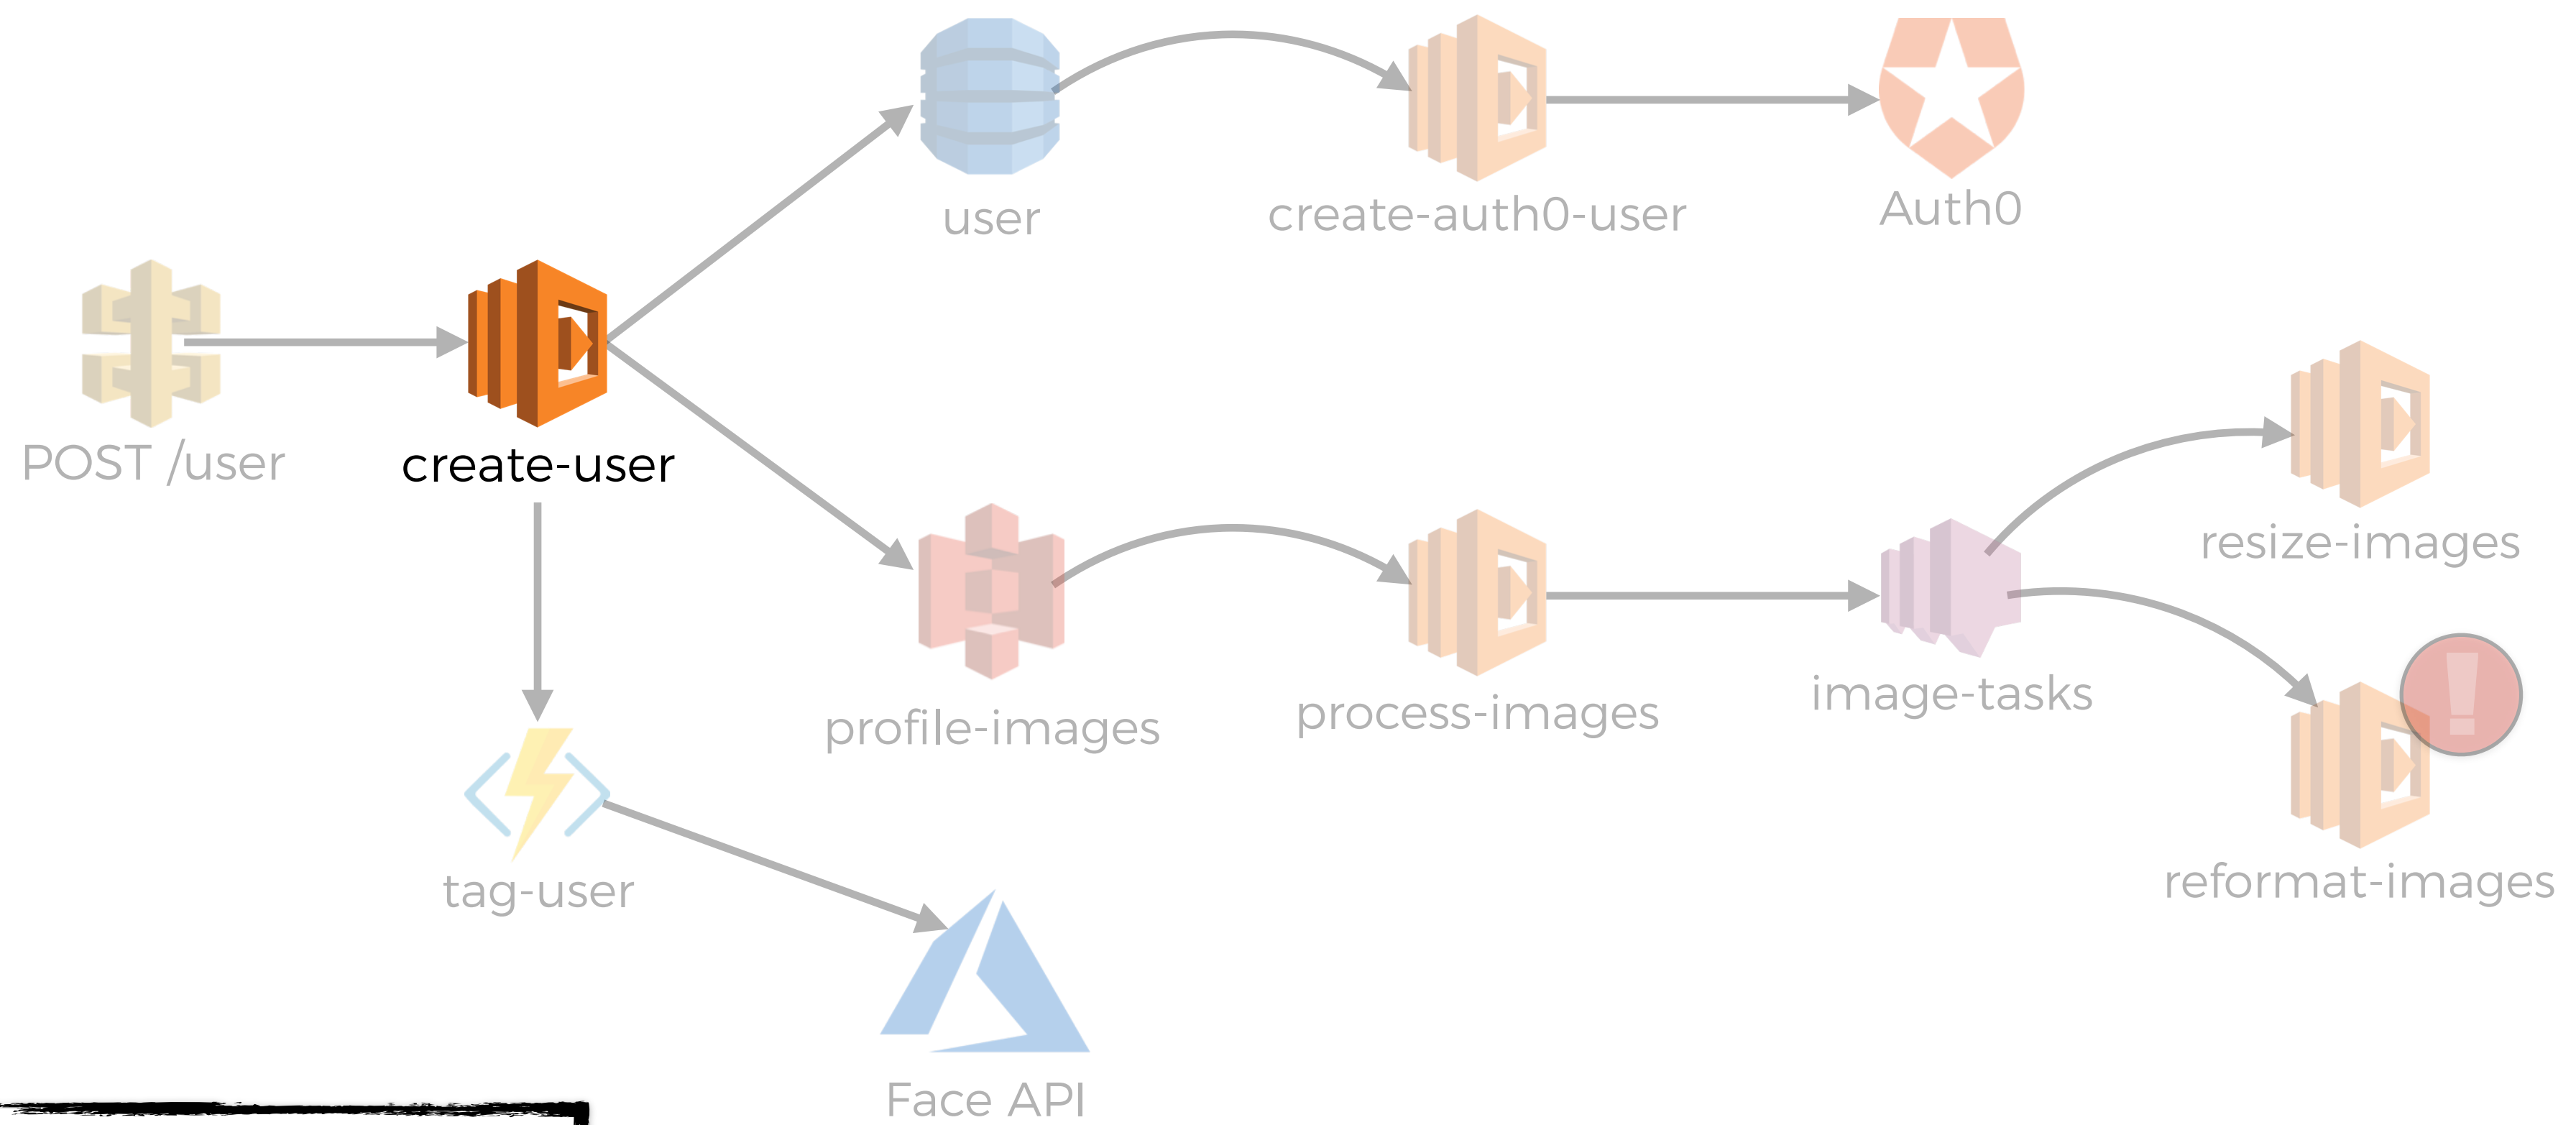
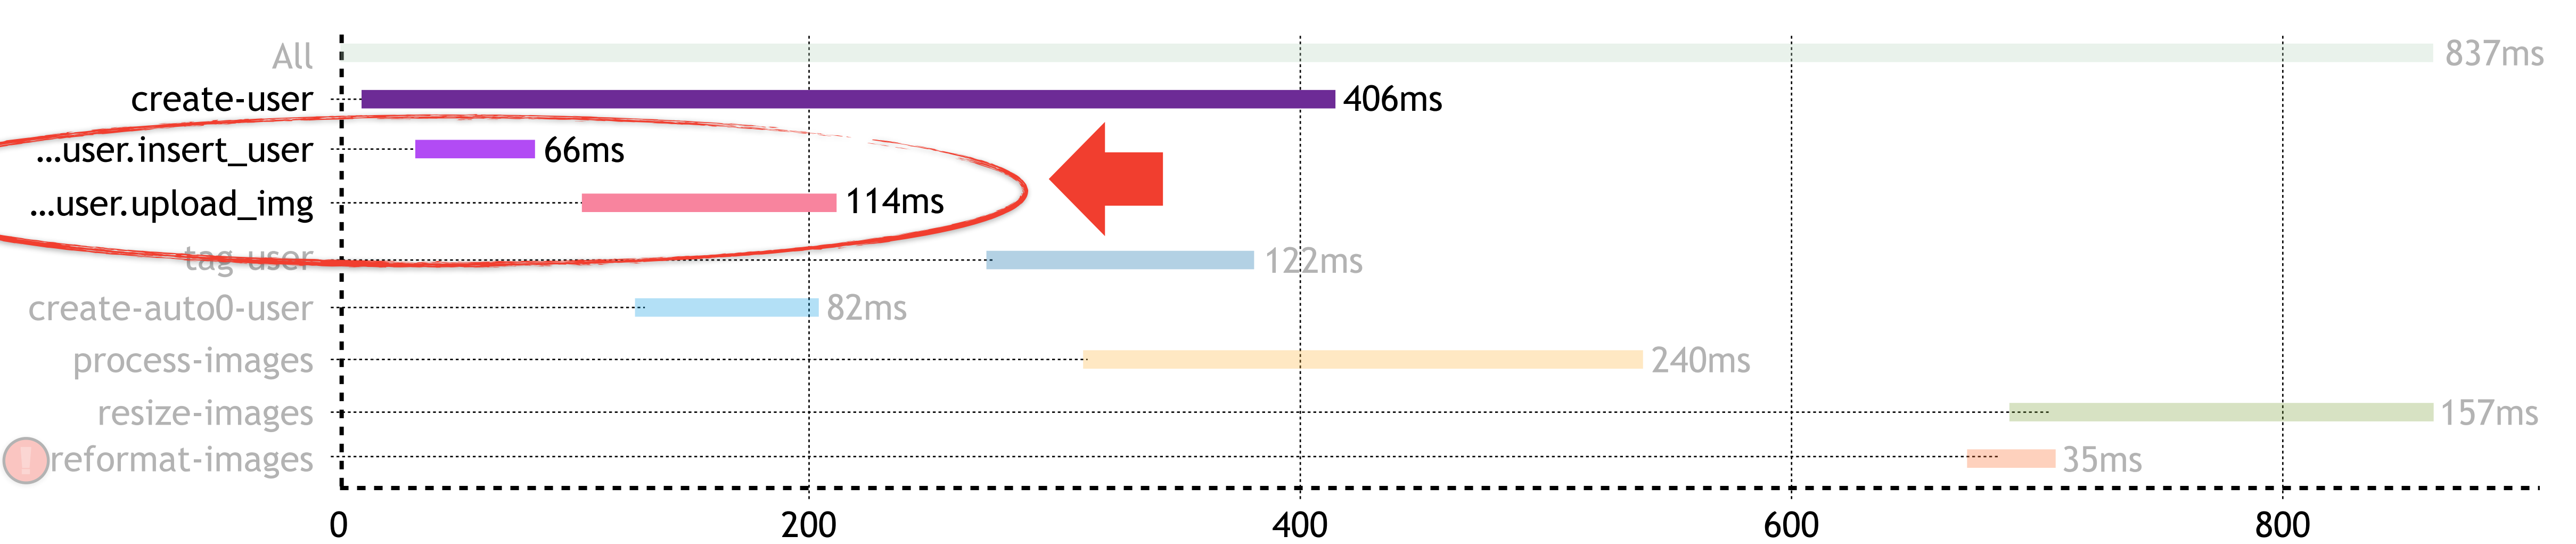
timestamp	component	level	message
-----------	-----------	-------	---------



Logs

Input/Output

timestamp	message	level	message
-----------	---------	-------	---------



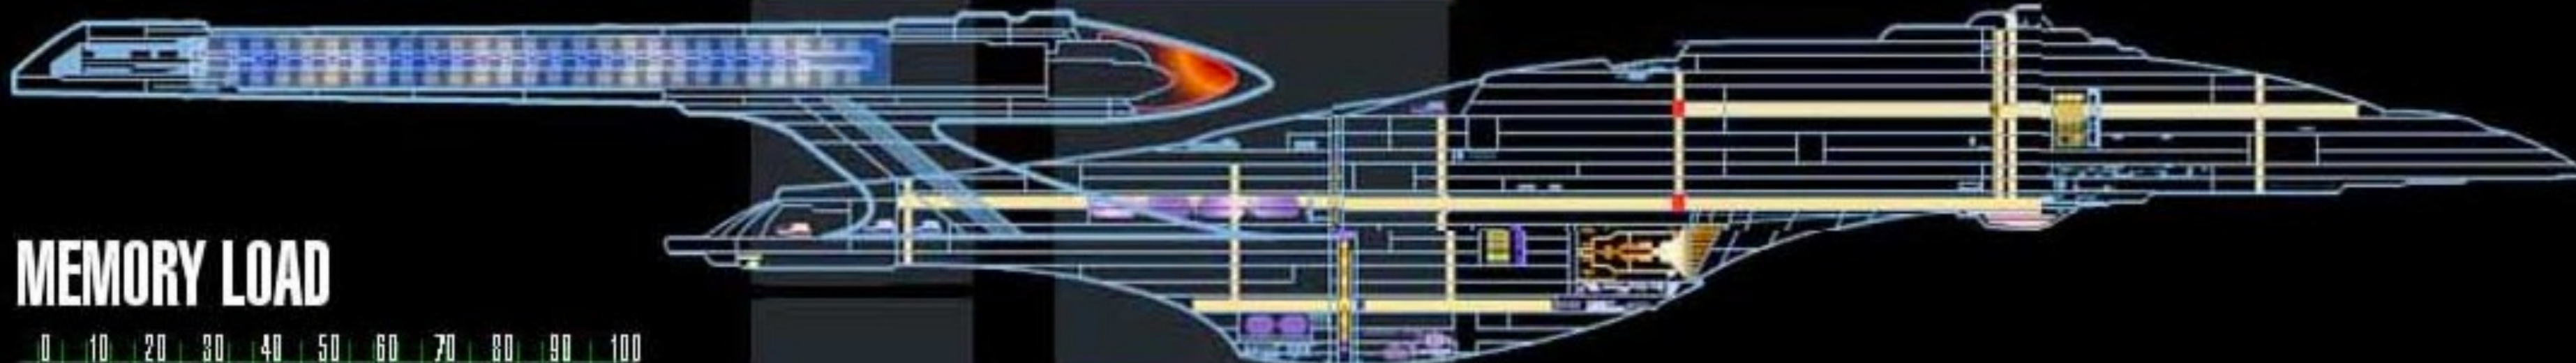
Logs **Input/Output**

timestamp message level message

all your **TRACING** needs in one place

CONTROL

168834	578988780787847	8788	
057342	84638385085003358	587483	4884387844
	47232308	4748	847379573
457732	5875755830038	4858883	28
83345	98087	472847	88783488447
			88
			88



- LEVEL 1 DIAGNOSTIC START
- LEVEL 2 DIAGNOSTIC START
- LEVEL 3 DIAGNOSTIC START
- LEVEL 4 DIAGNOSTIC START
- LEVEL 5 DIAGNOSTIC START

LEVEL 3 DIAGNOSTIC - AUTO SYSTEM CHECKS WITH MINOR CREW VERIFICATION OVER TEN MINUTES. SYSTEMS ON-LINE.

MEMORY LOAD



AMD Athlon 1200 10%

Hard Drives:



- EXIT
- TOOLS
- WINDOWS

SYSTEM ONLINE

CRITICAL SYSTEMS REPORT

- LIFE SUPPORT STATUS: NORMAL
- SHIELDS STATUS: NORMAL
- WEAPONS STATUS: NORMAL
- WARP CORE STATUS: NORMAL

LONG-RANGE SCANNERS STATUS: NORMAL
RADIATION LEVELS: LOW

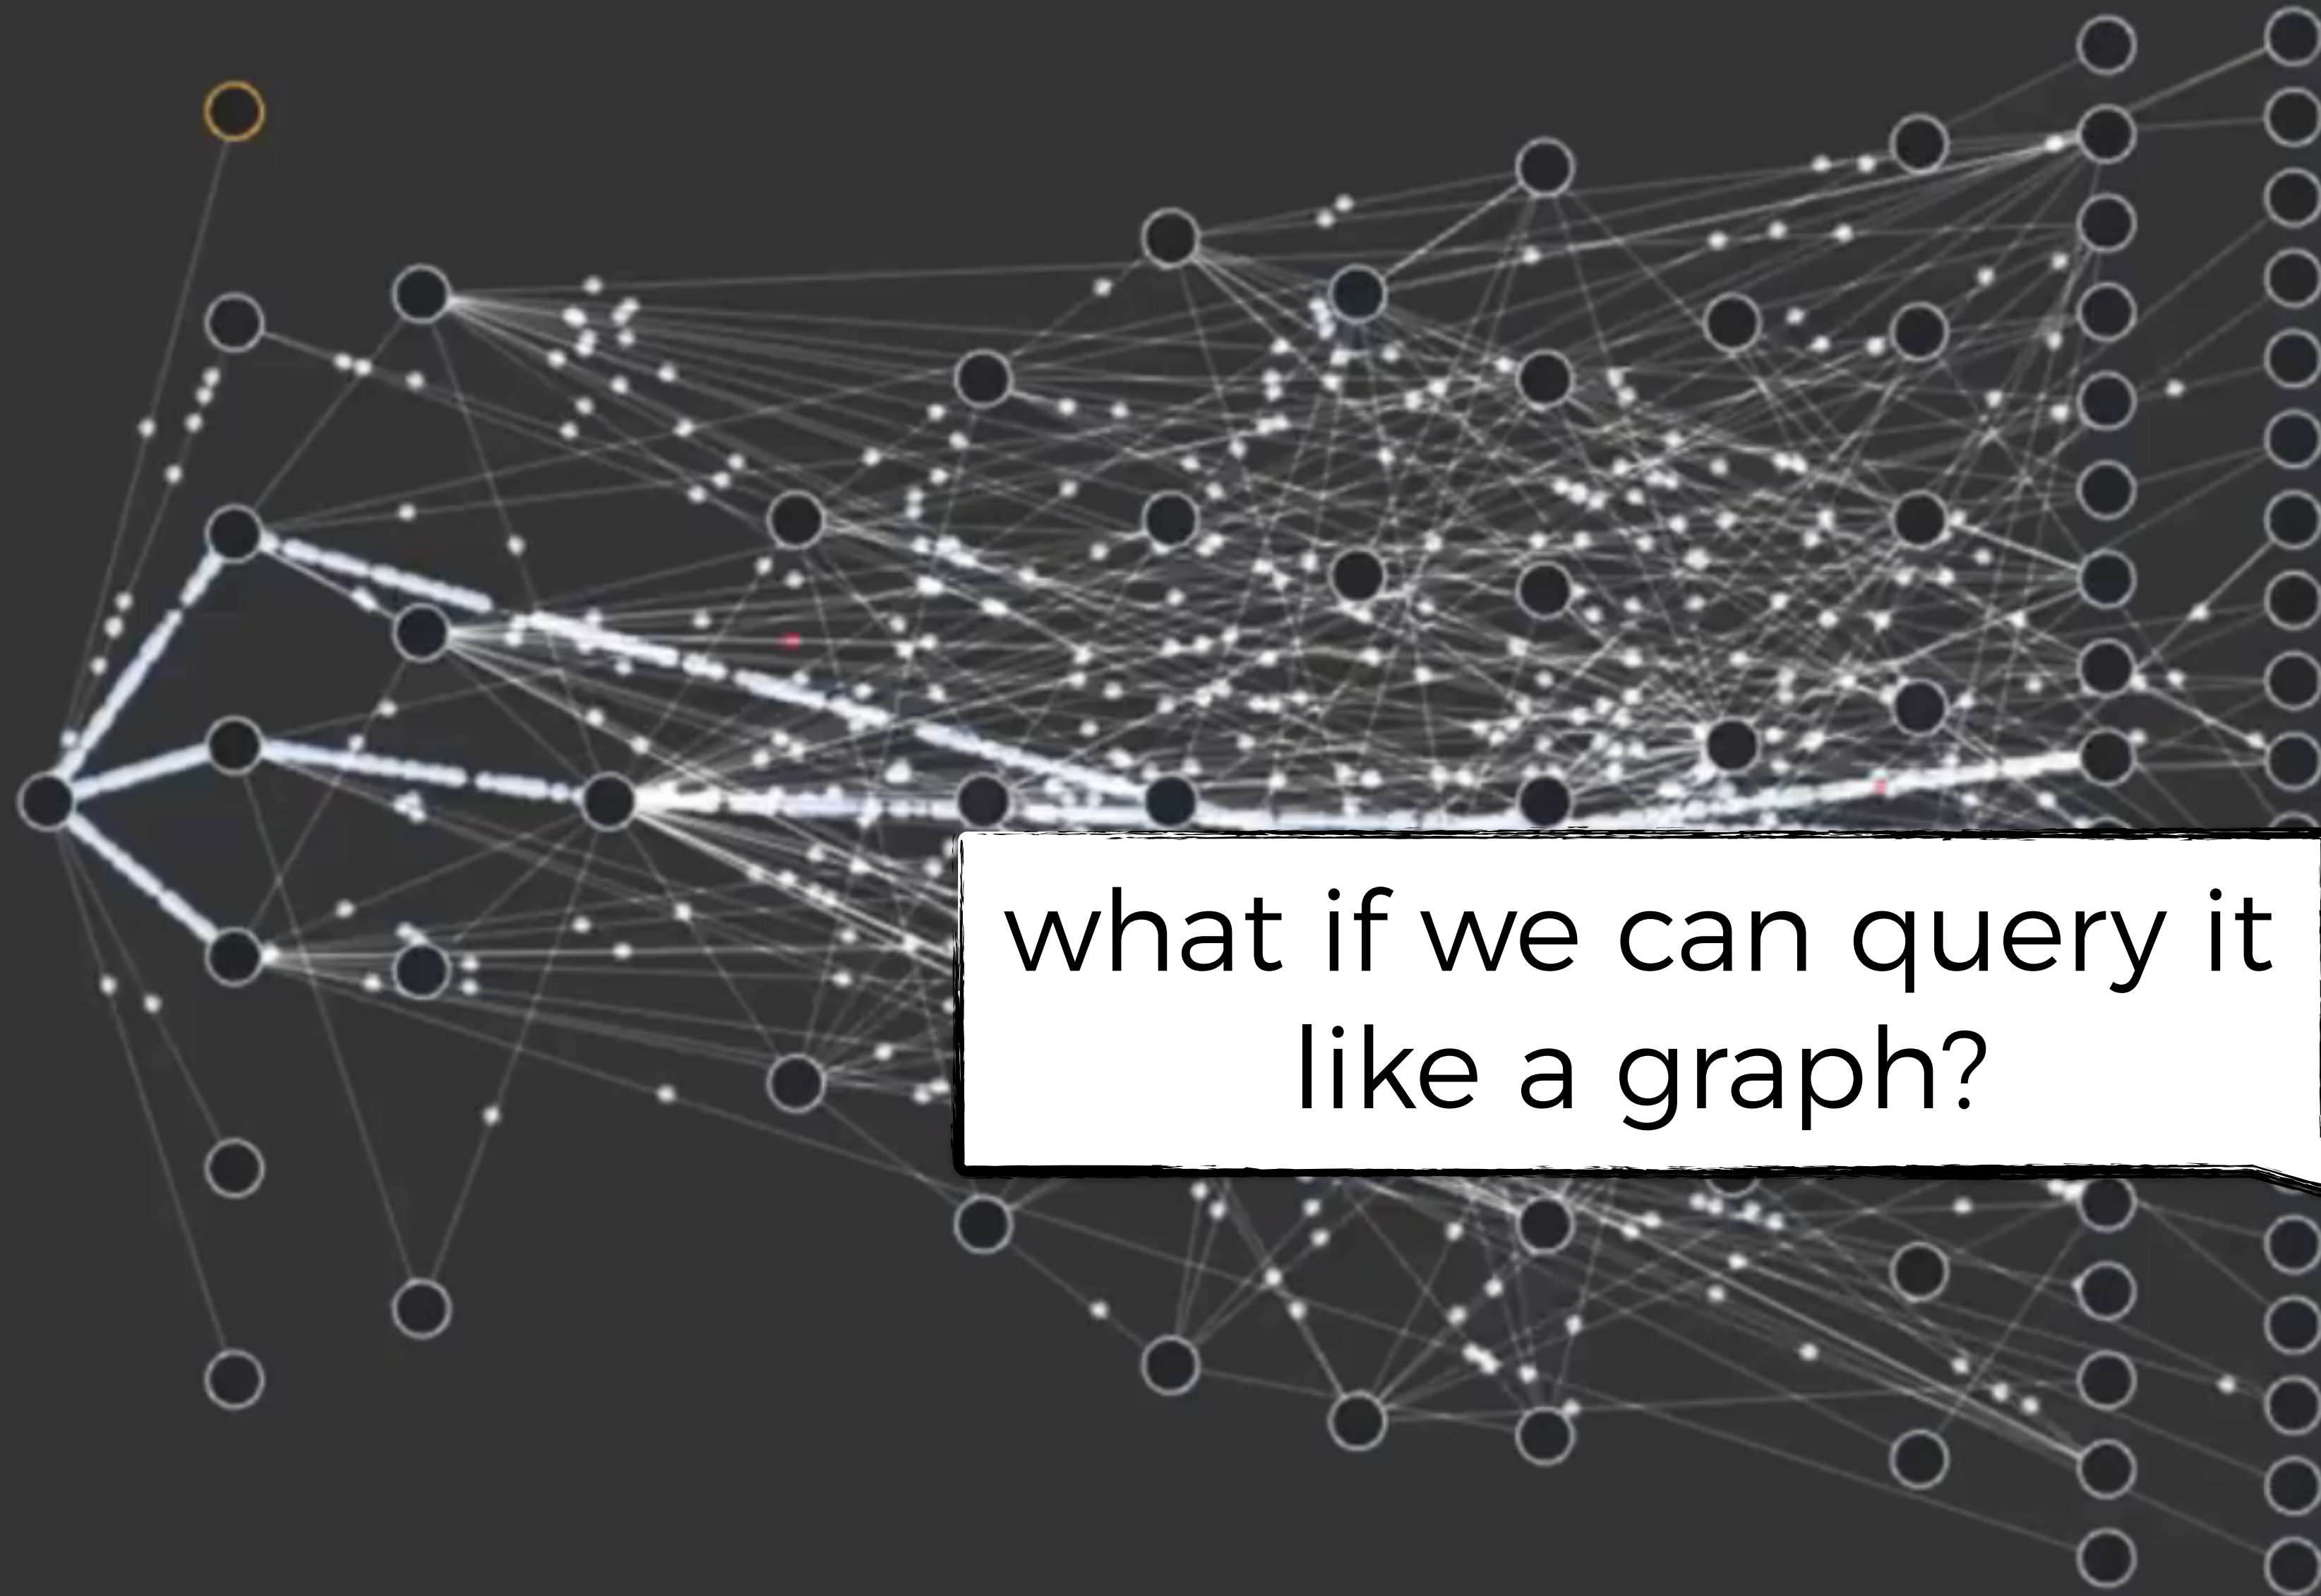






mmm... it's a graph





what if we can query it like a graph?





Amazon Neptune – A Fully Managed Graph Database Service

by Randall Hunt | on 29 NOV 2017 | in [Amazon Neptune](#), [Amazon RDS*](#), [AWS Re:Invent*](#), [Database*](#), [Launch*](#), [News*](#) | [Permalink](#) | [Comments](#) | [Share](#)

Of all the data structures and algorithms we use to enable our modern lives, **graphs** are changing the world everyday. Businesses continuously create and ingest rich data with complex relationships. Yet developers are still forced to model these complex relationships in traditional databases. This leads to frustratingly complex queries with high costs and increasingly poor performance as you add relationships. We want to make it easy for you to deal with these modern and increasingly complex datasets, relationships, and patterns.

Hello, Amazon Neptune

Today we're launching a limited preview ([sign up here](#)) of [Amazon Neptune](#), a fast and reliable graph database service that makes it easy to gain insights from relationships among your highly connected datasets. The core of Amazon Neptune is a purpose-built, high-performance graph database engine optimized for storing billions of relationships and querying the graph with milliseconds of latency. Delivered as a fully managed database, Amazon Neptune frees customers to focus on their applications rather than tedious undifferentiated operations like maintenance, patching, backups, and restores. The service supports fast-failover, point-in-time recovery, and Multi-AZ deployments for high availability. With support for up to 15 read replicas you can scale query throughput to 100s of thousands of queries per second. Amazon Neptune runs within your [Amazon Virtual Private Cloud](#) and allows you to encrypt your data at rest, giving you complete control over your data integrity in transit and at rest.

DATABASE SERVICES

Amazon Neptune

Fast, reliable graph database

Amazon Neptune is a fast, reliable graph database service that makes it easy to build and run applications that work with highly connected datasets.

Graph Database

Preview of Amazon Neptune

<http://amzn.to/2nk7uiW>

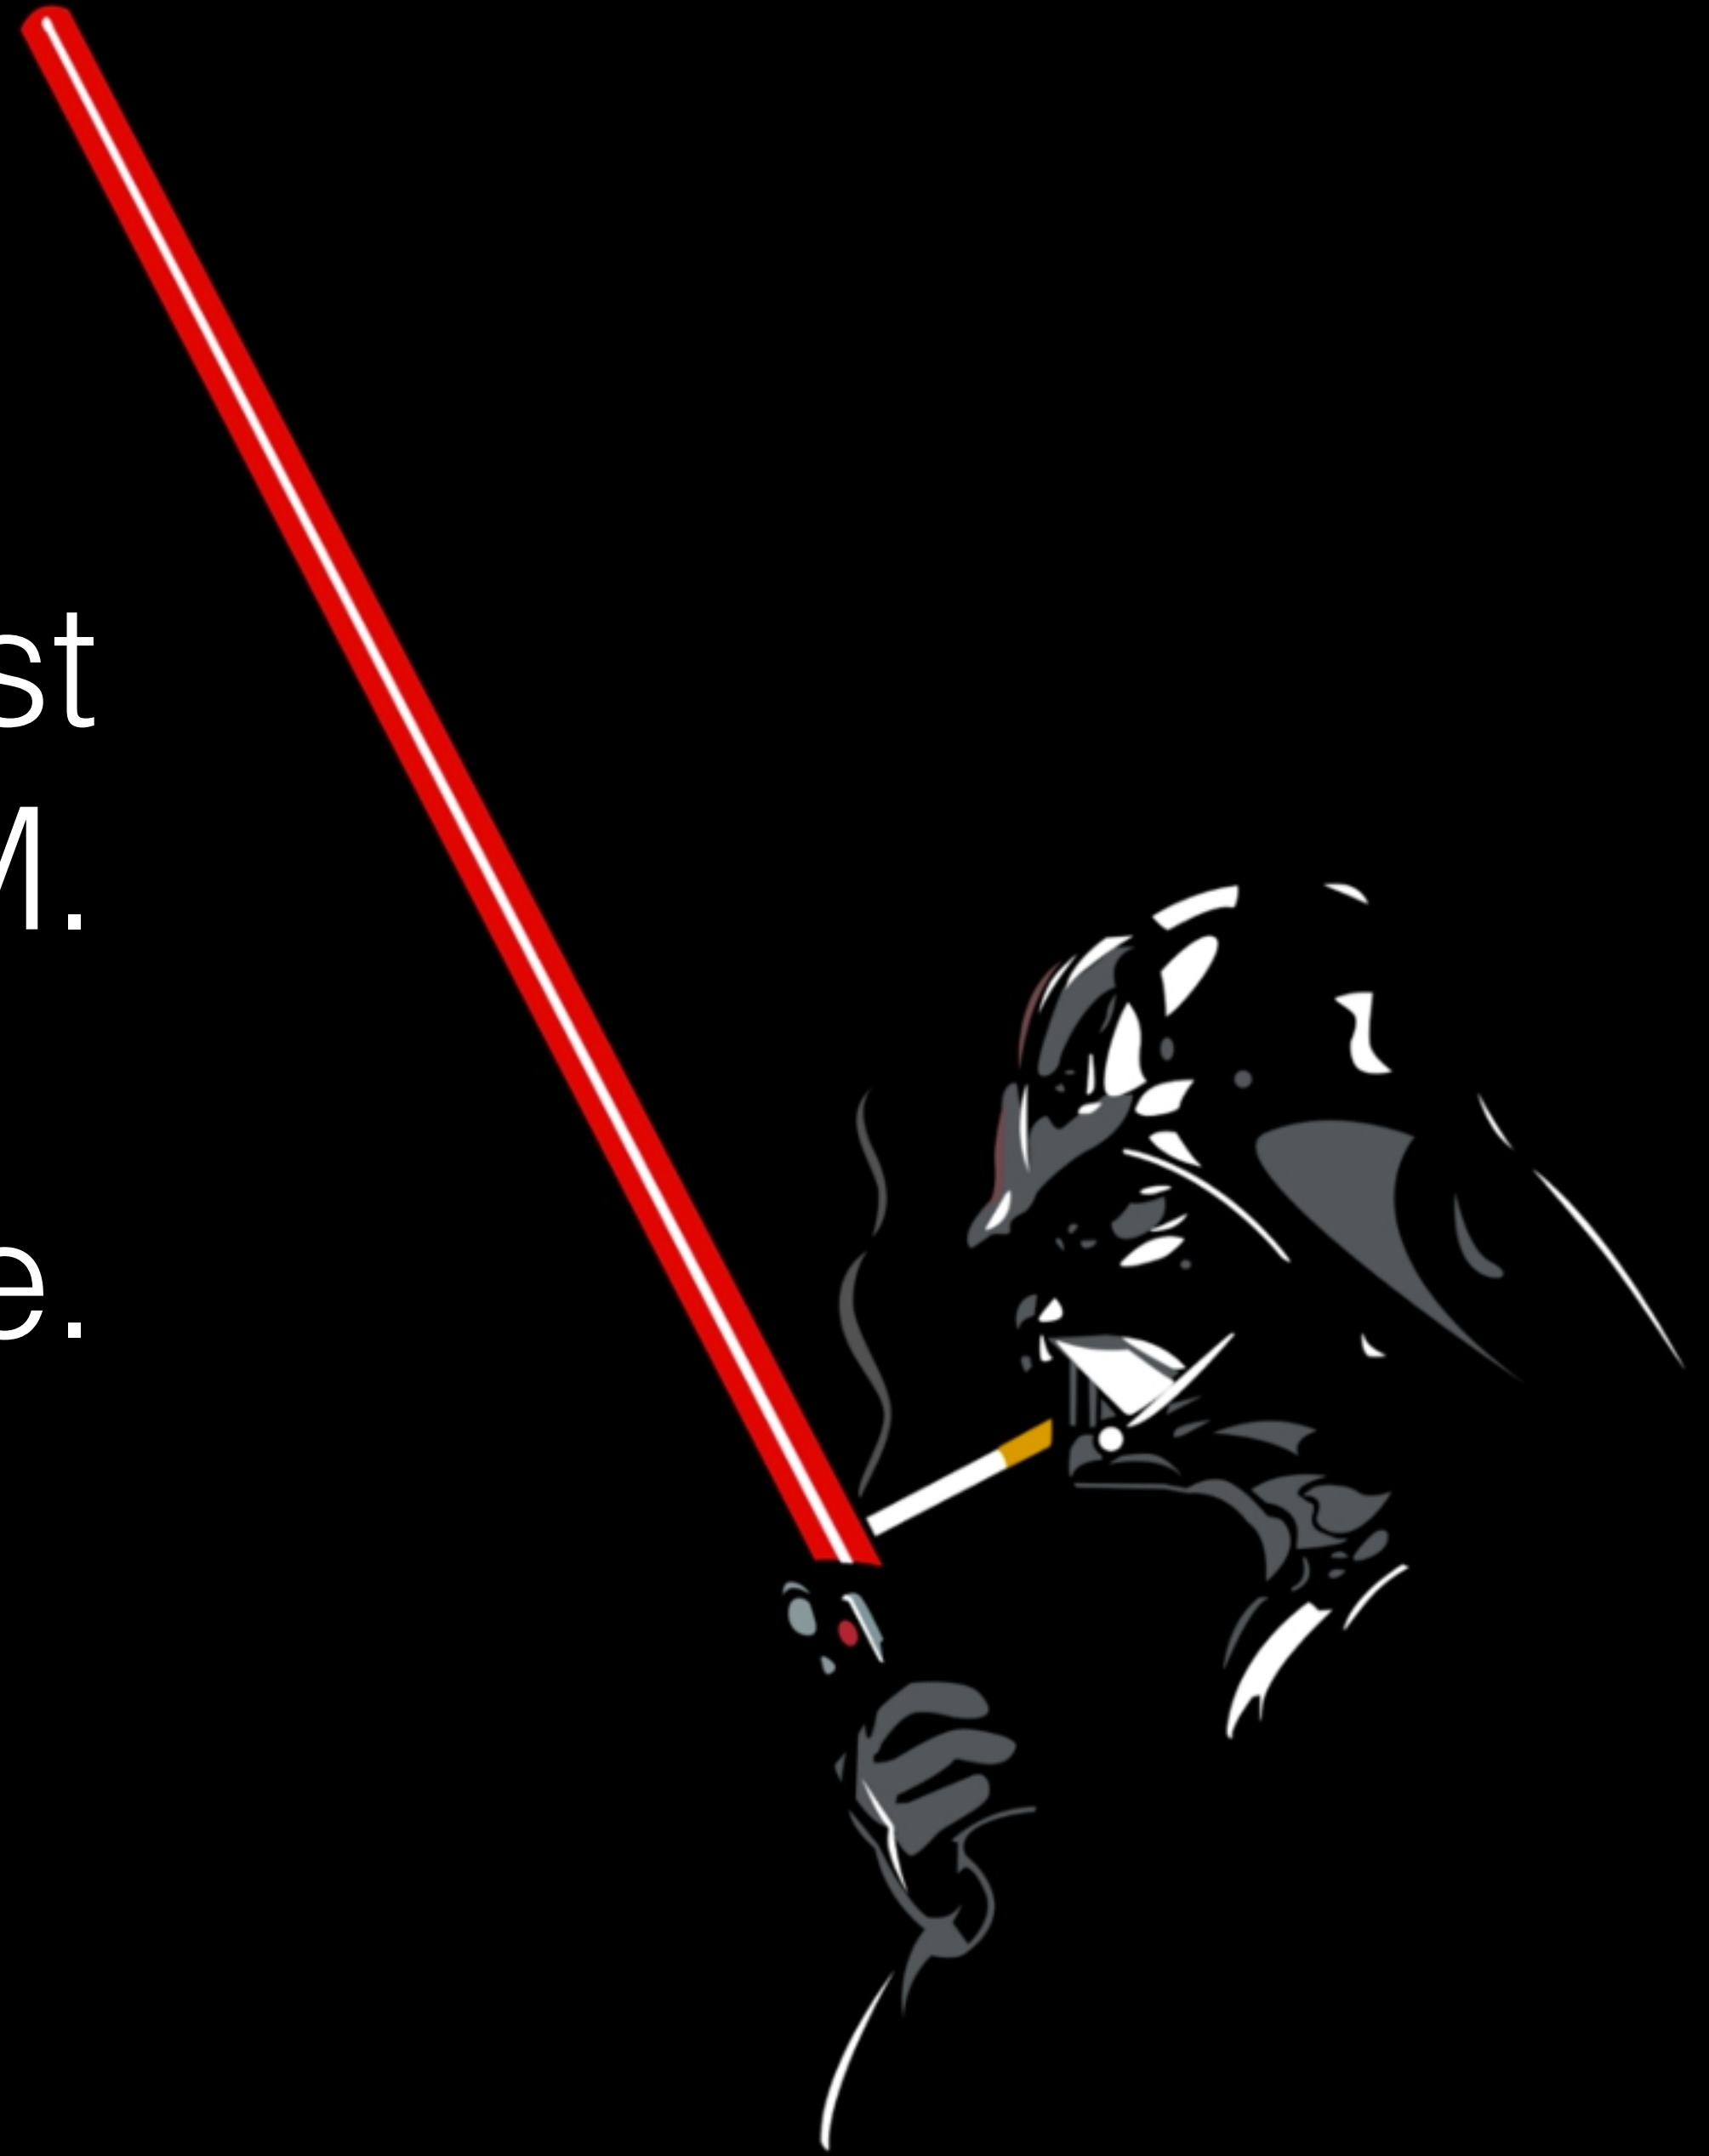
ability to query based on the **relationship**
between observed components
(as well as the components themselves)

root **cause** analysis

“ *the elevated error rate in service X was **caused** by DynamoDB table throttling.* ”

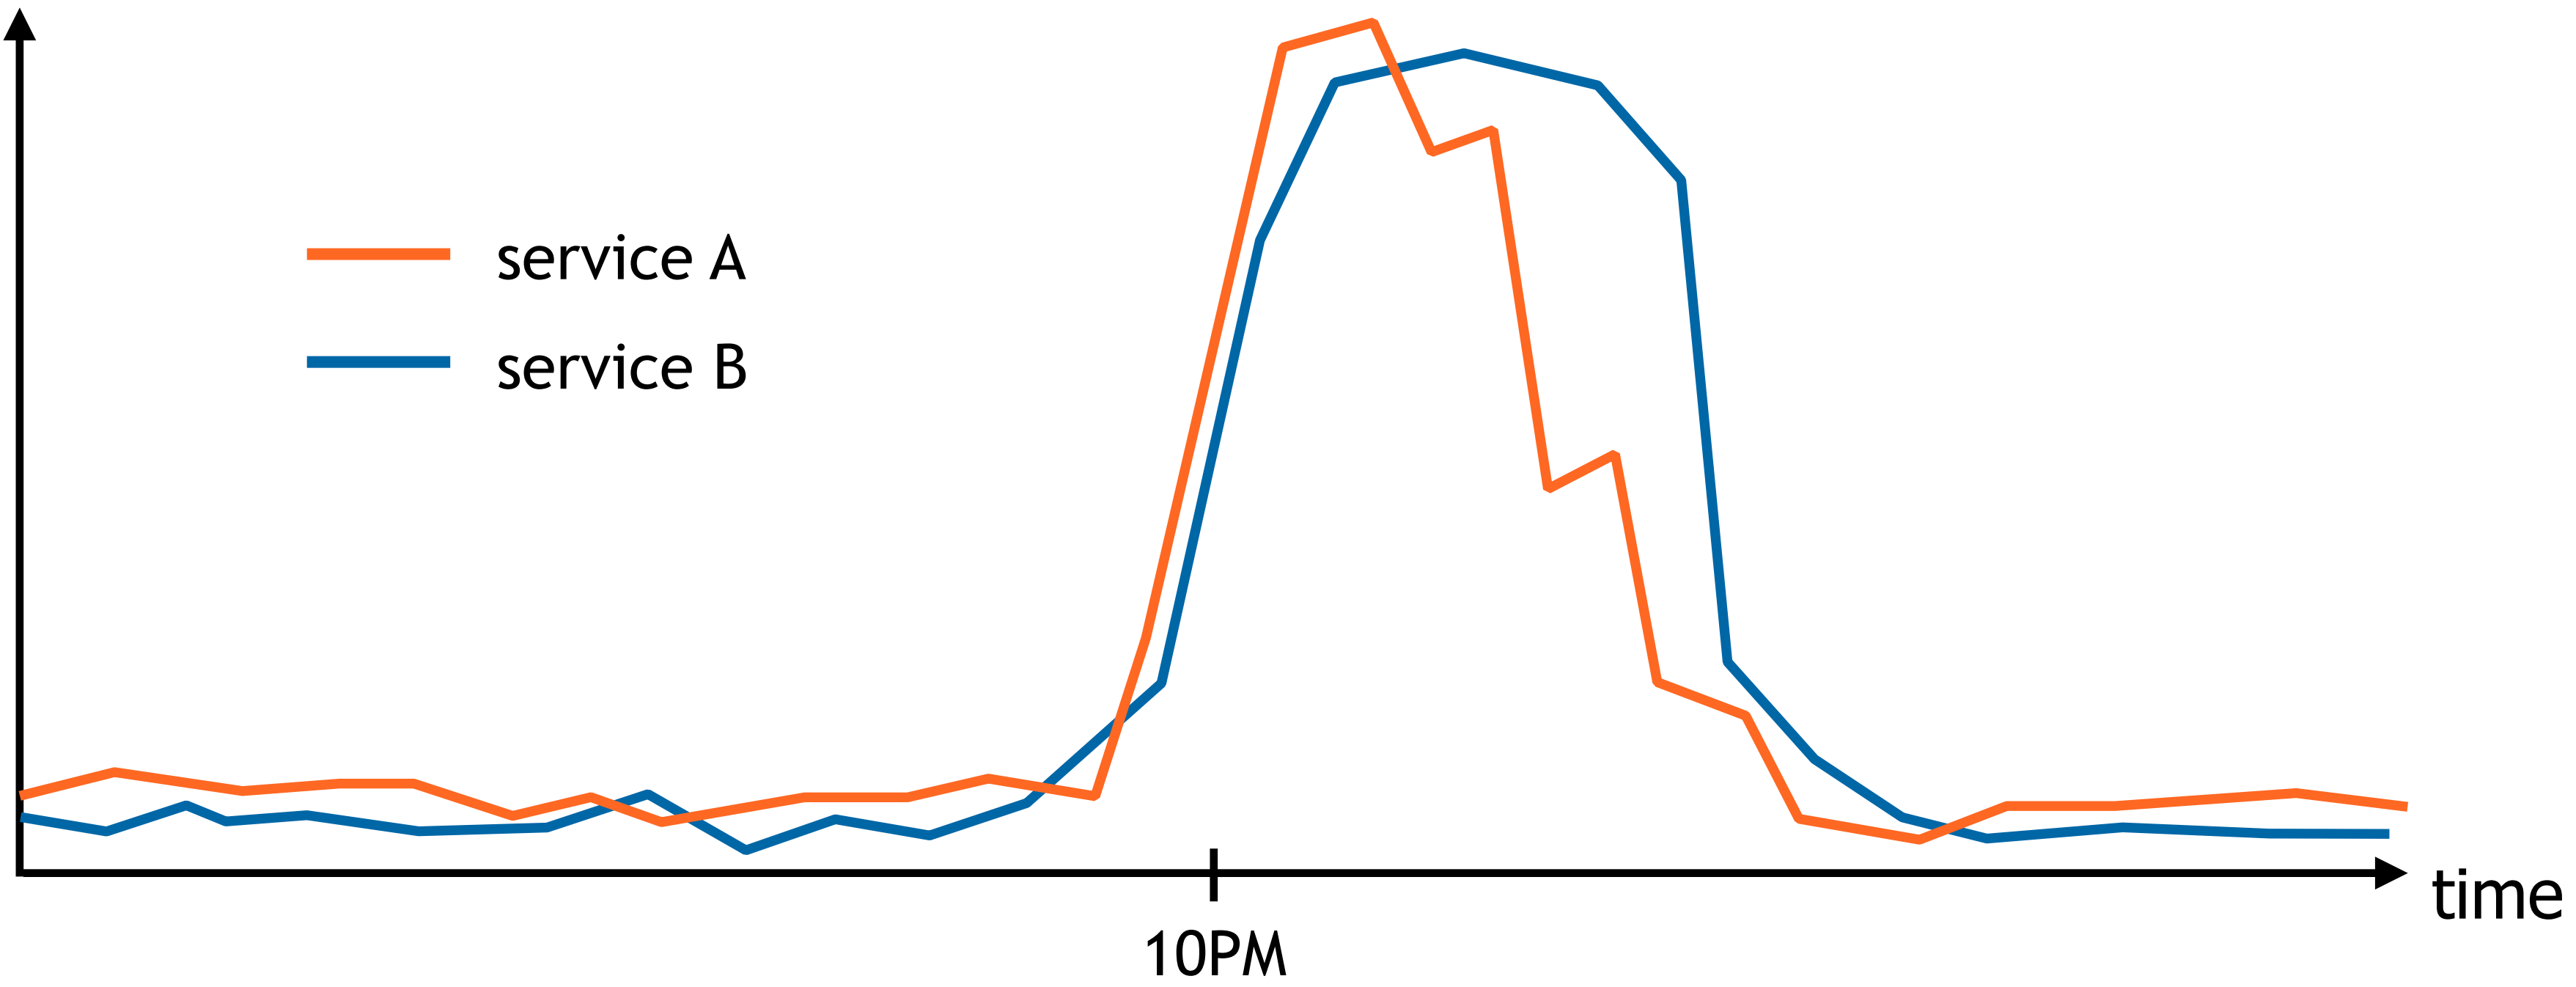
payment was slow last
night around 10PM.

investigate.



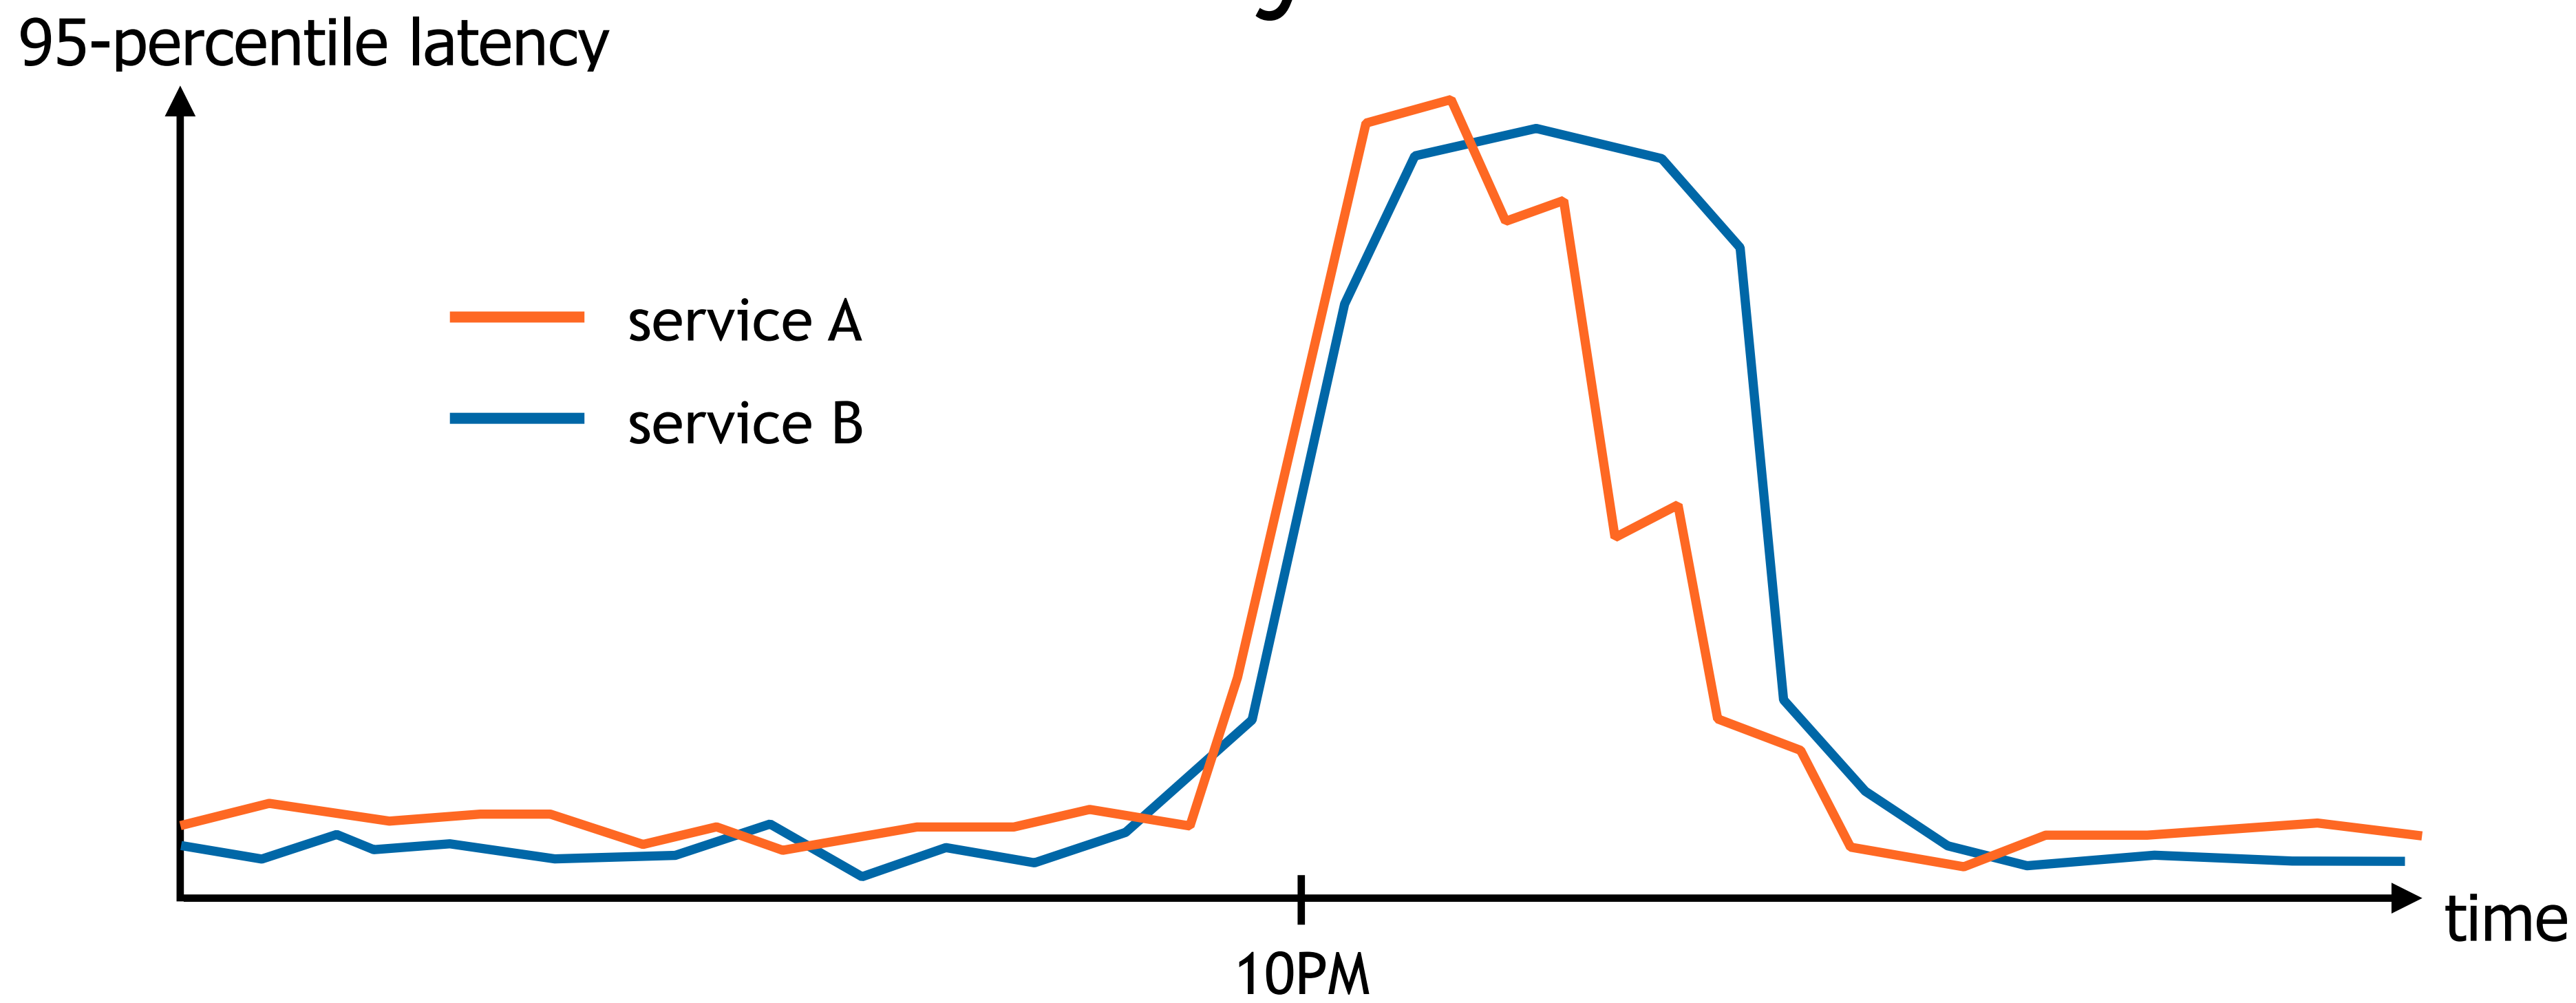


95-percentile latency

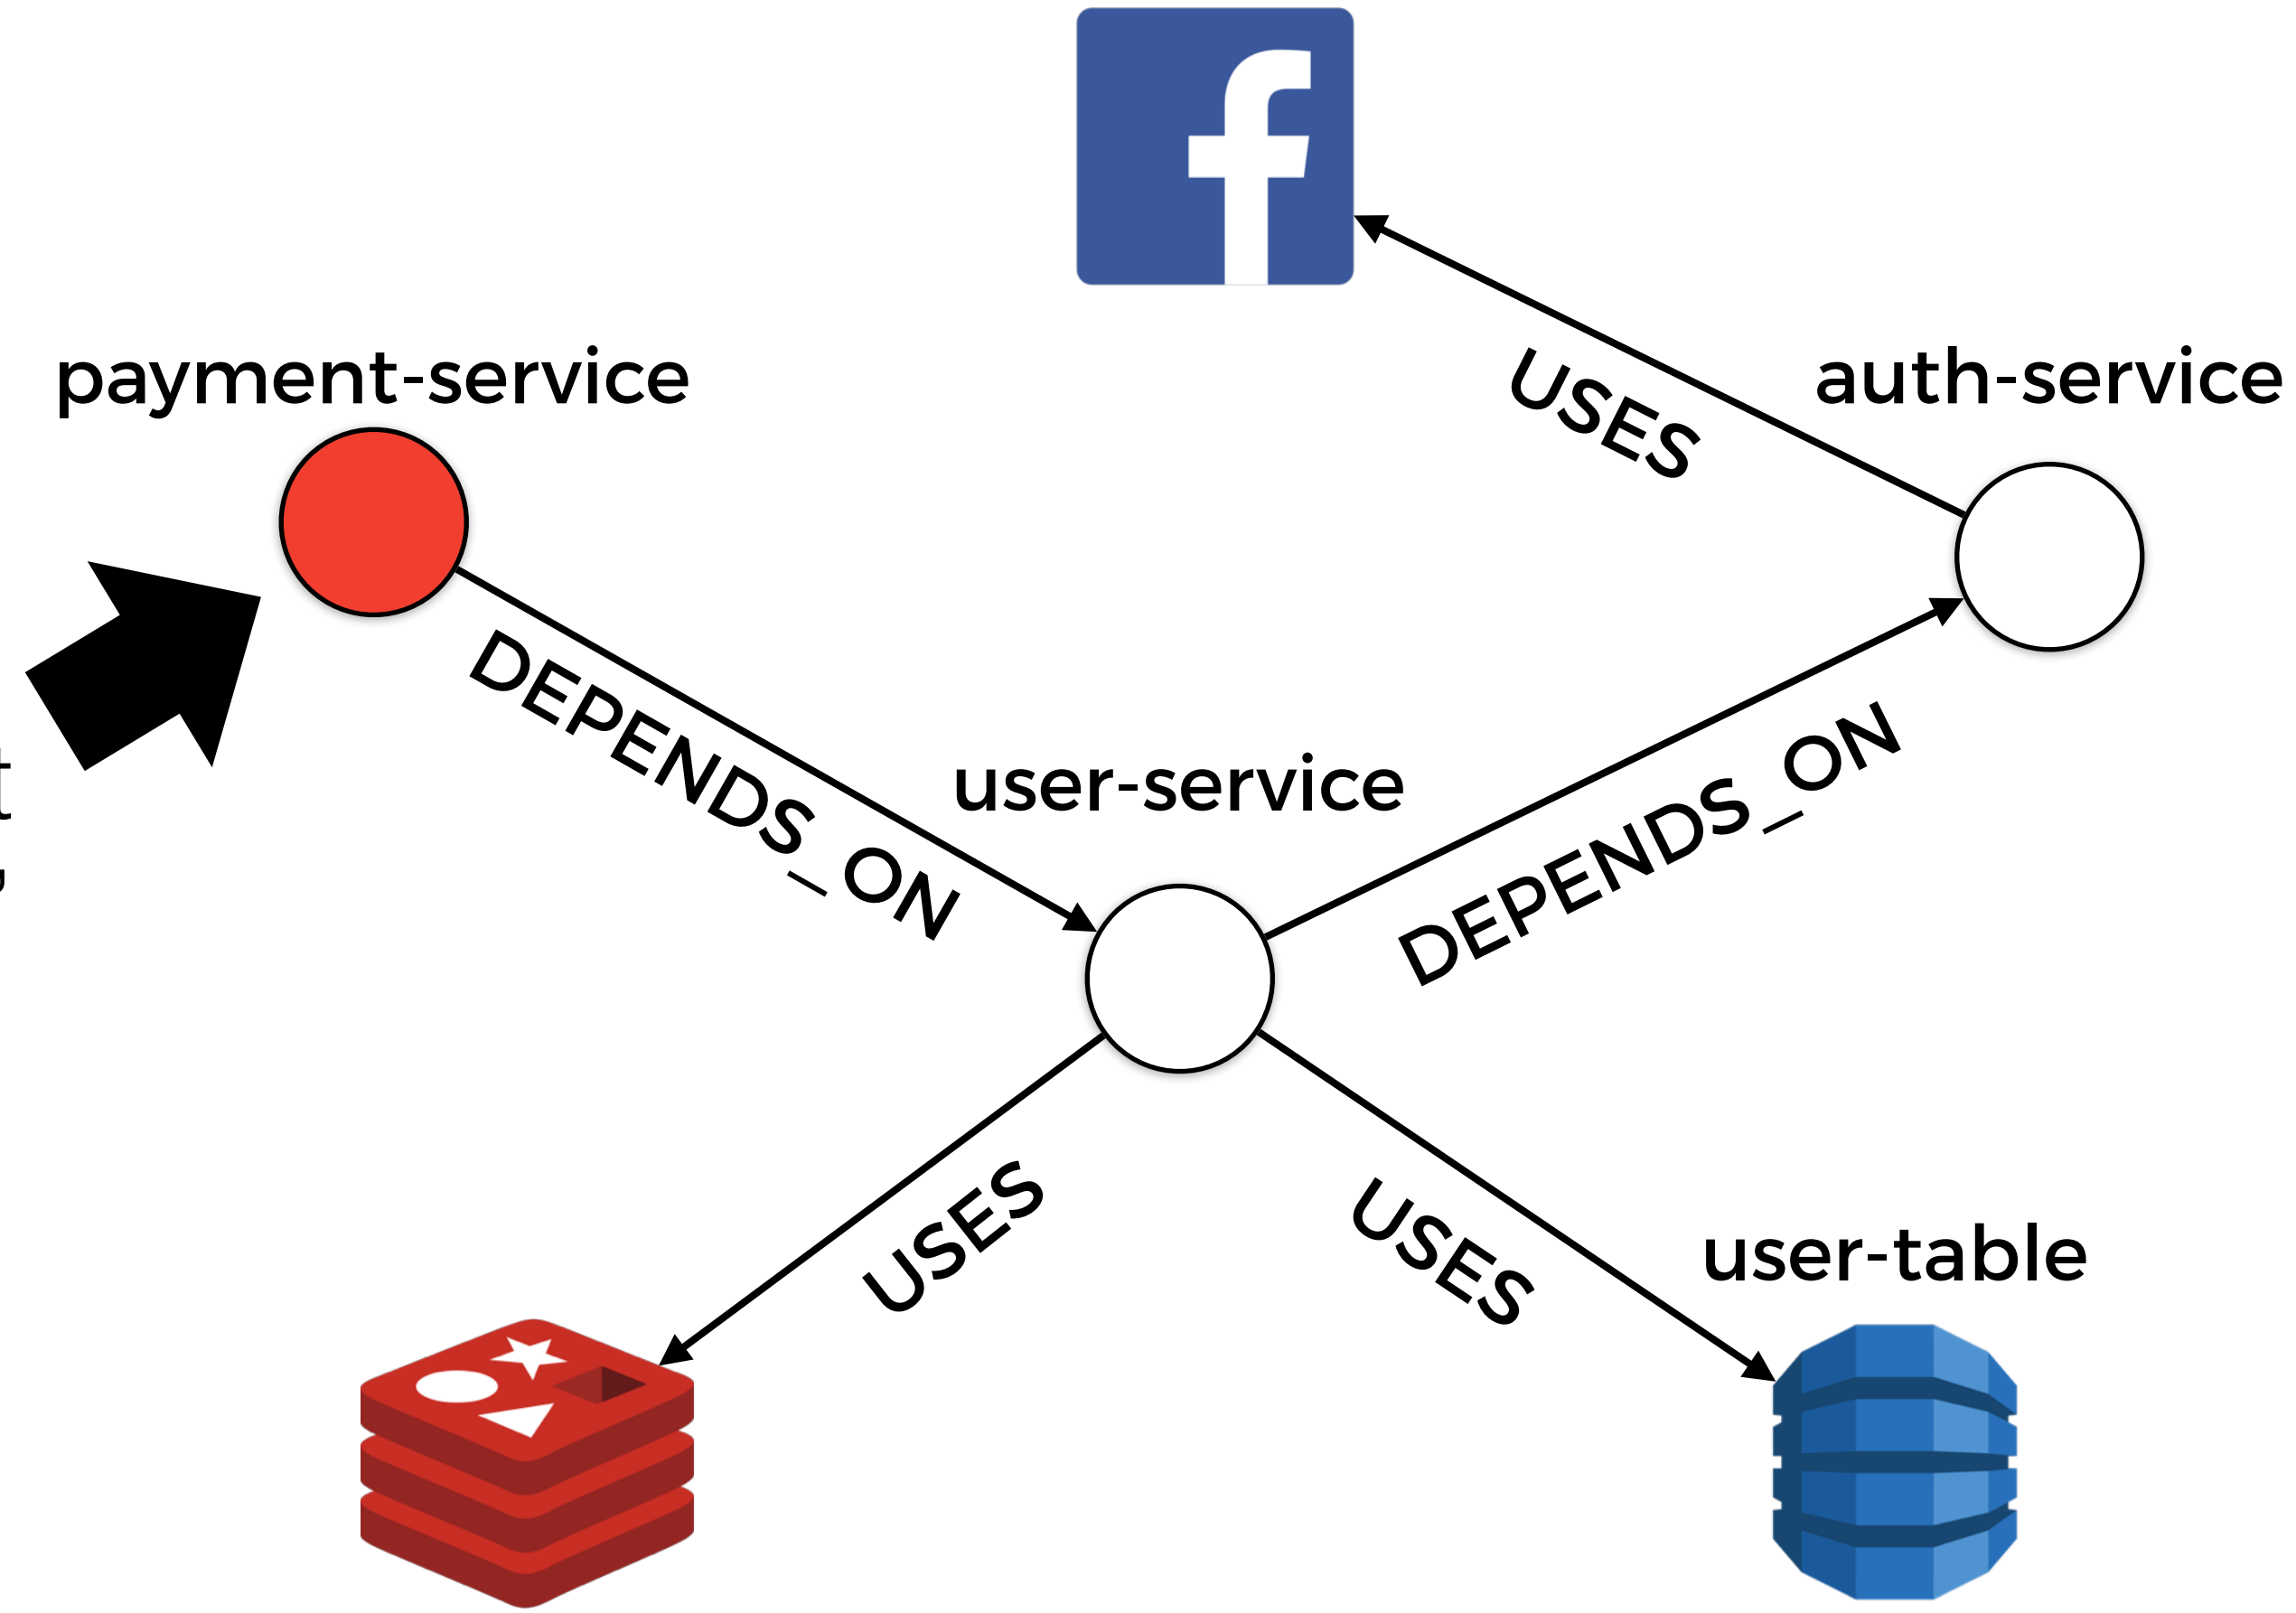


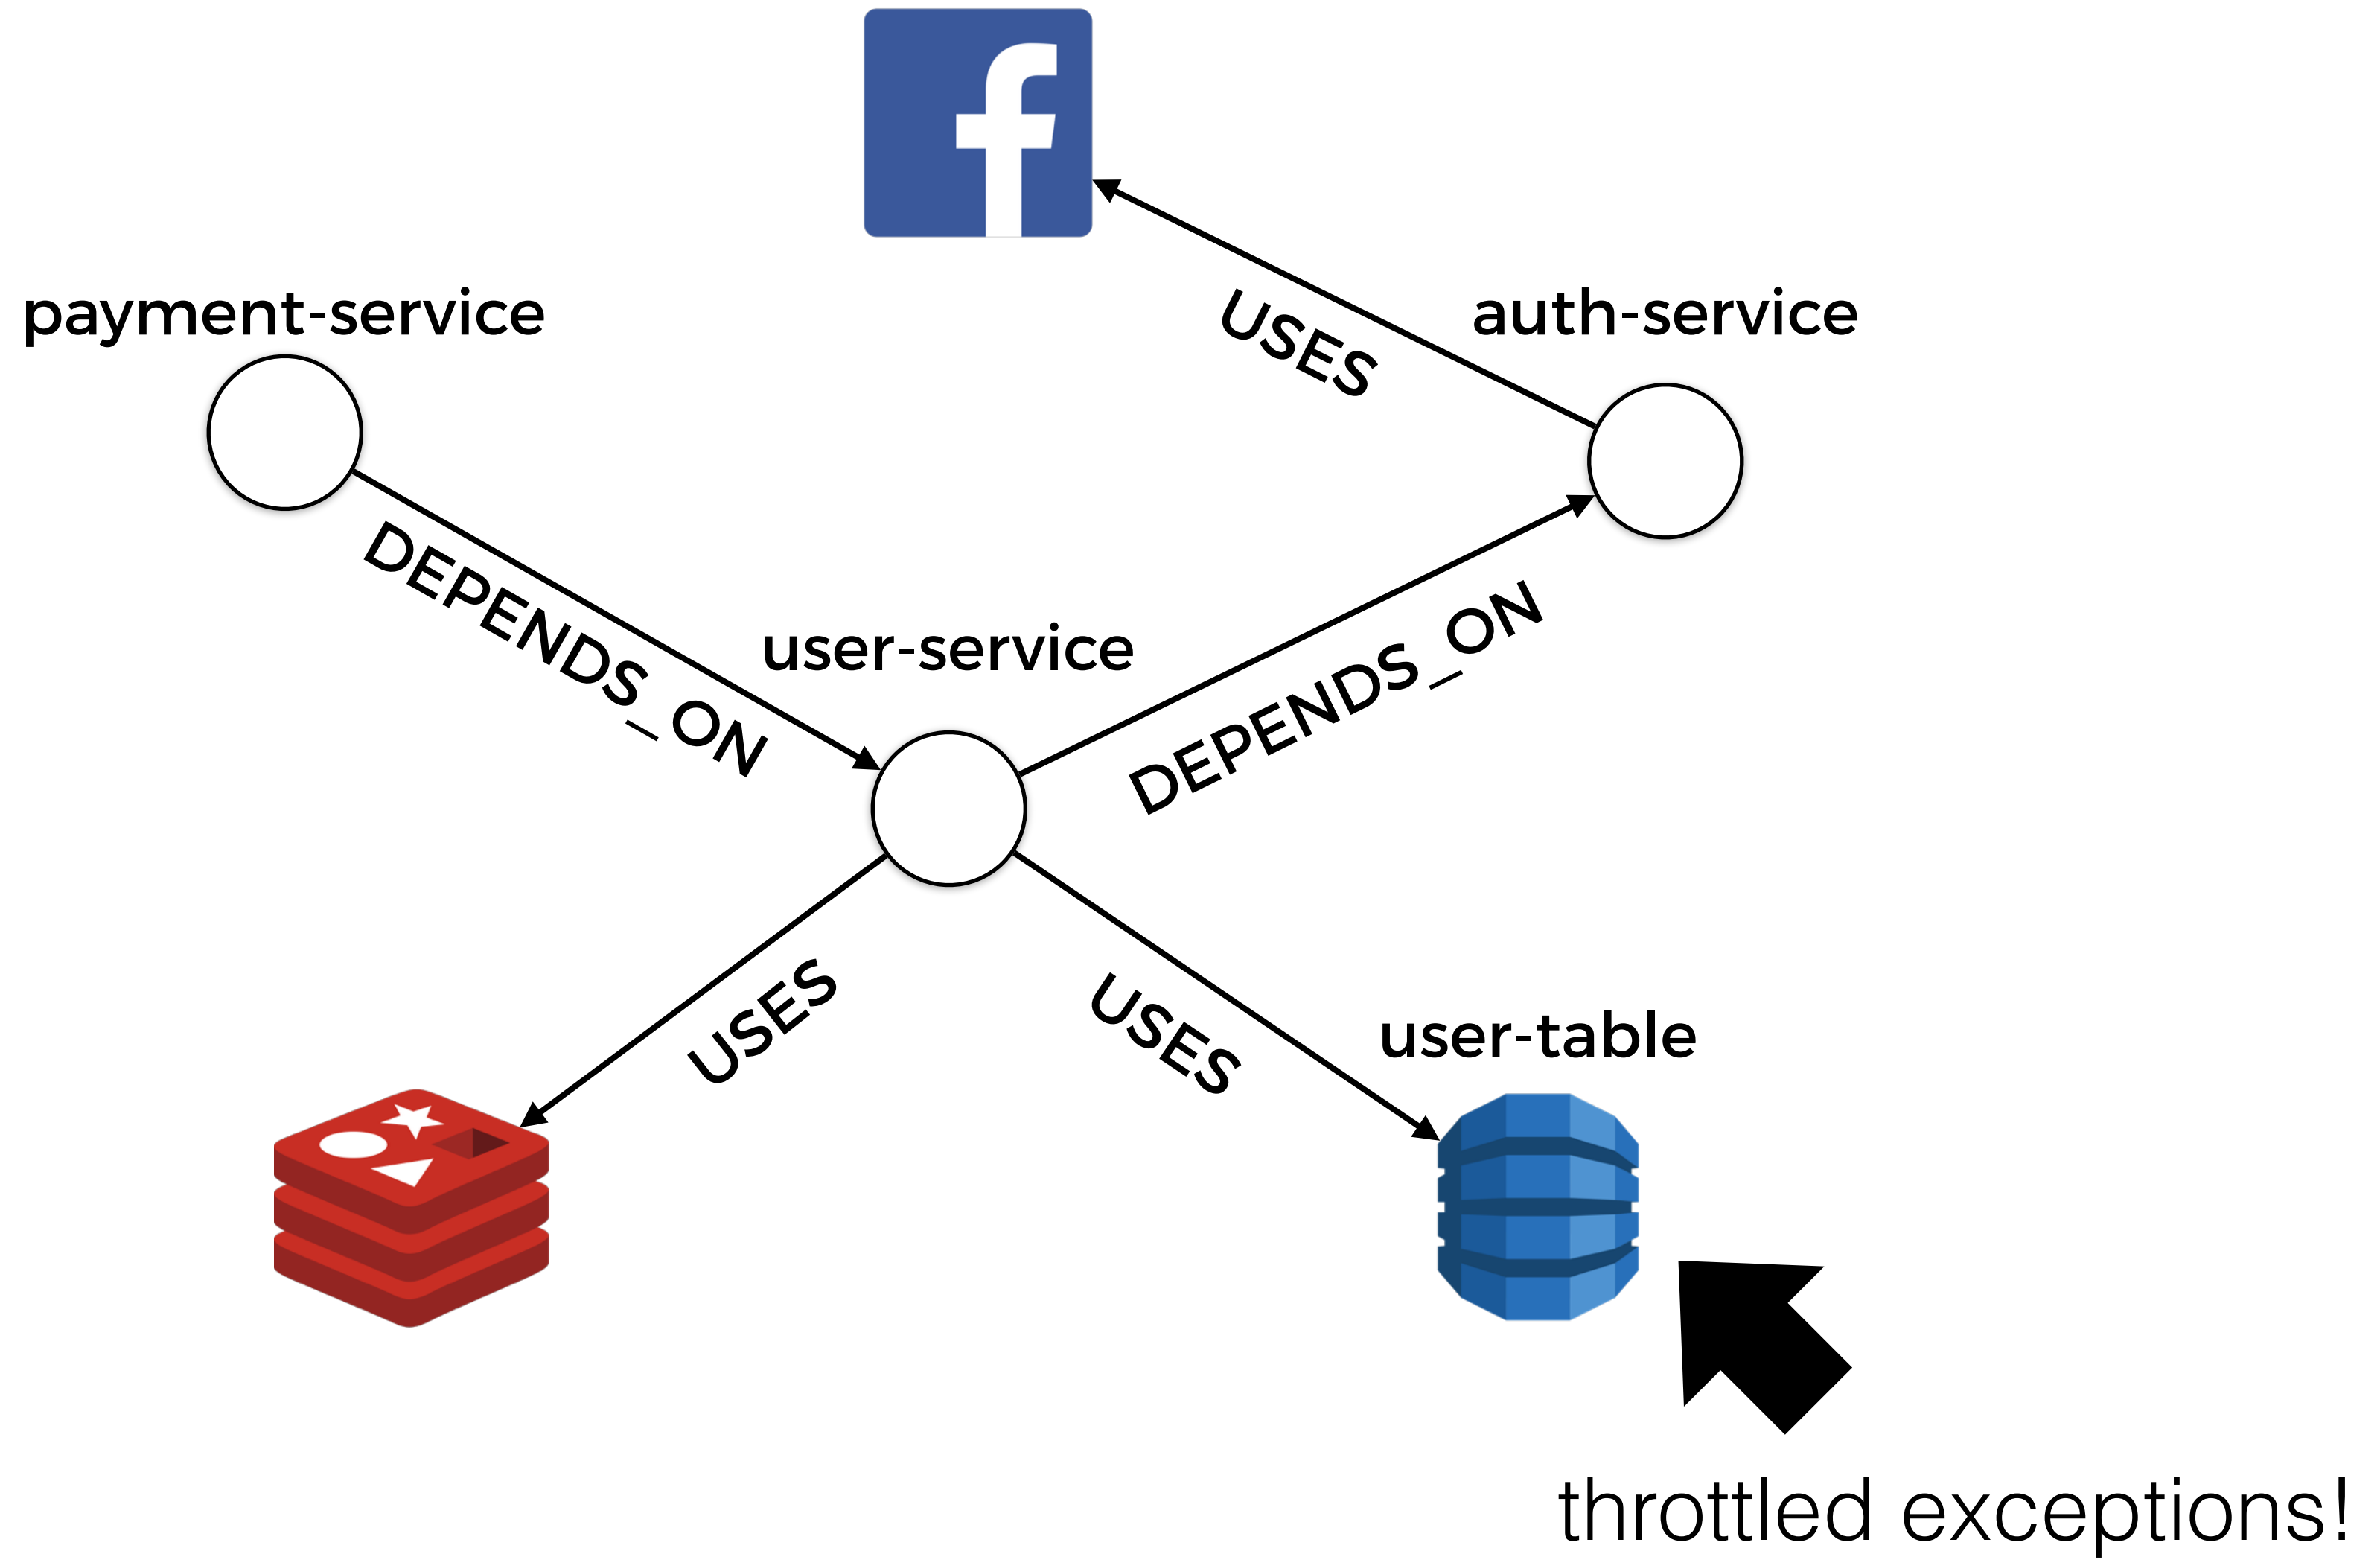


causality? or correlation?

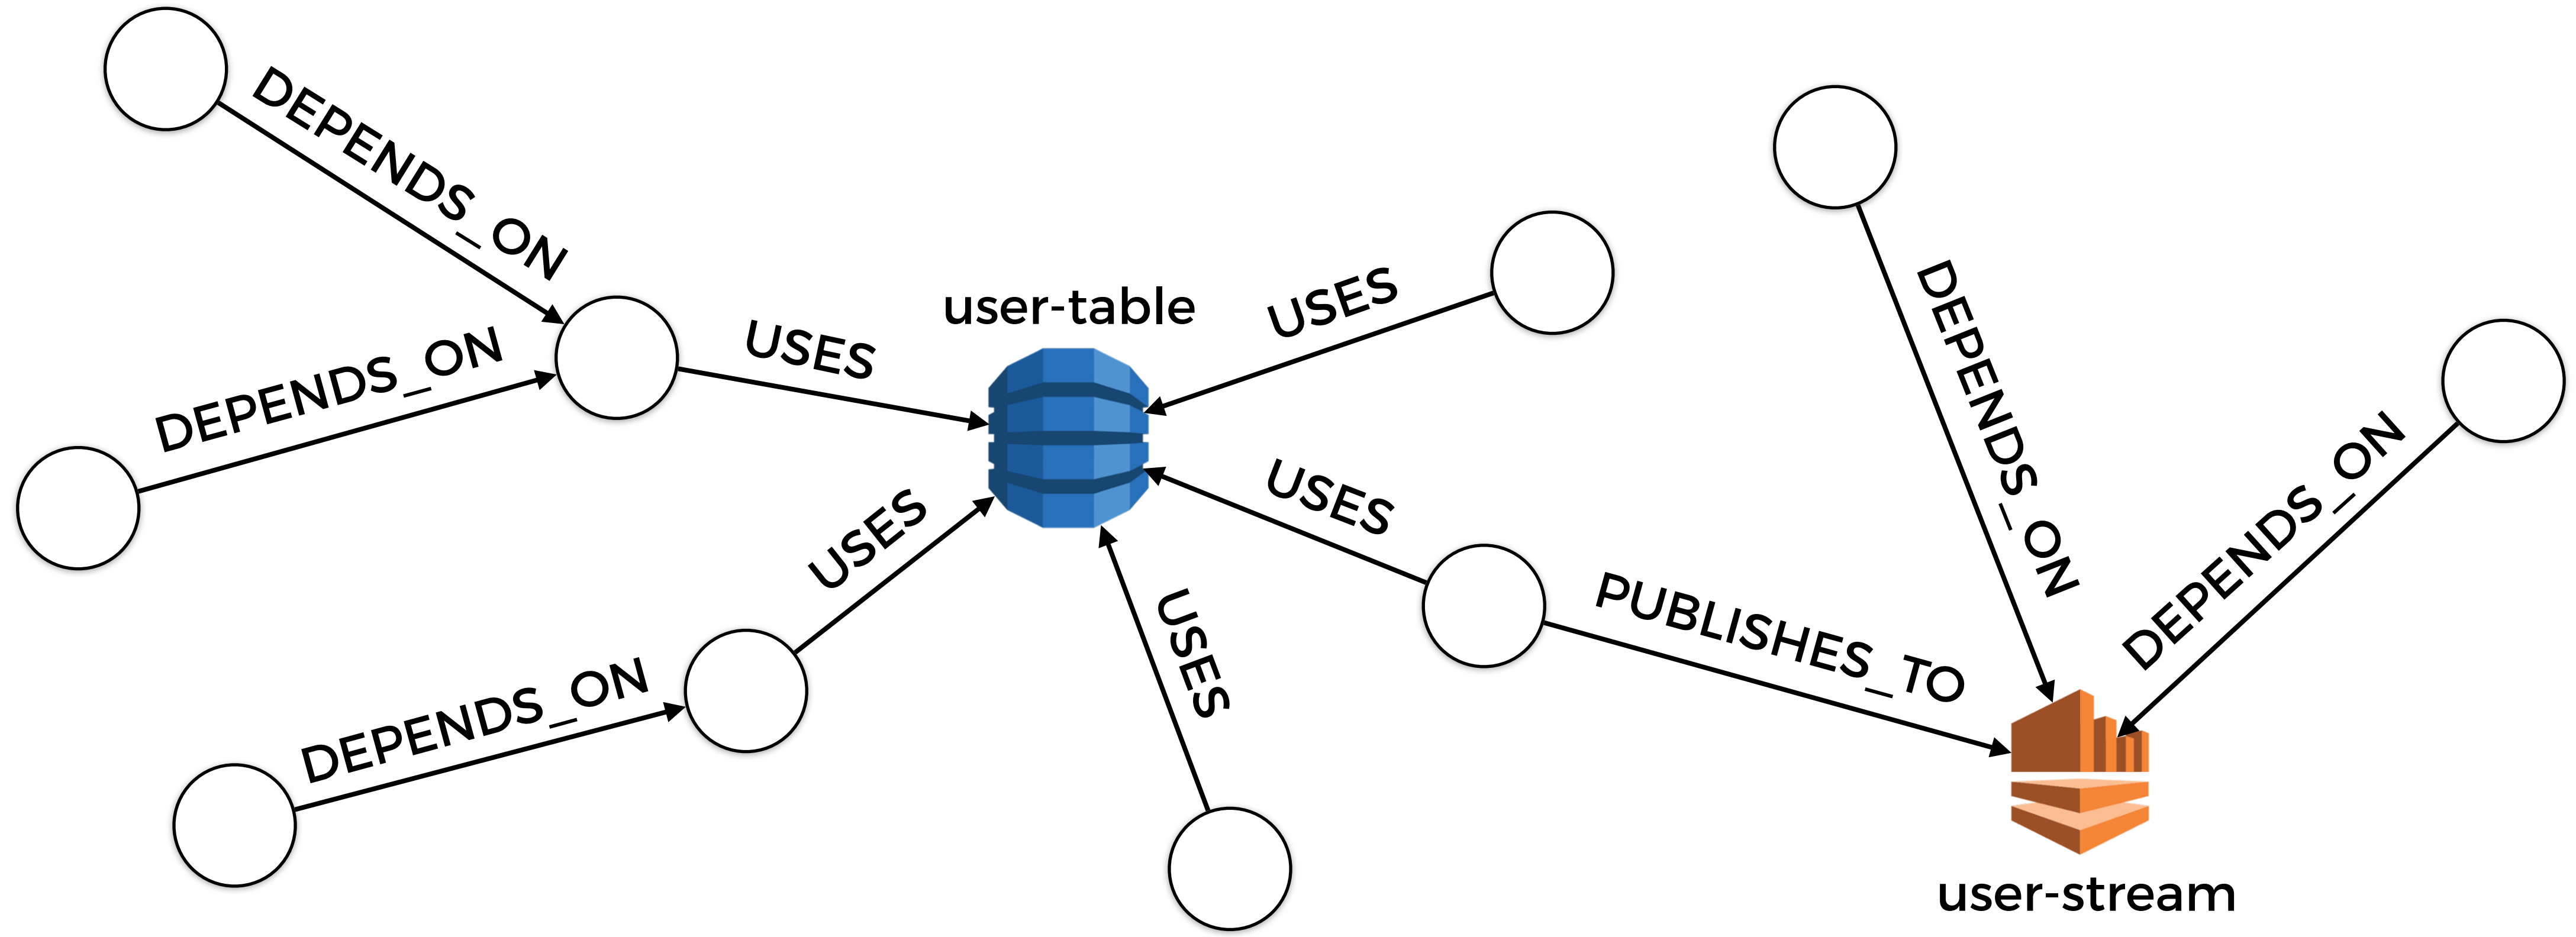


“payment was slow last night around 10PM”

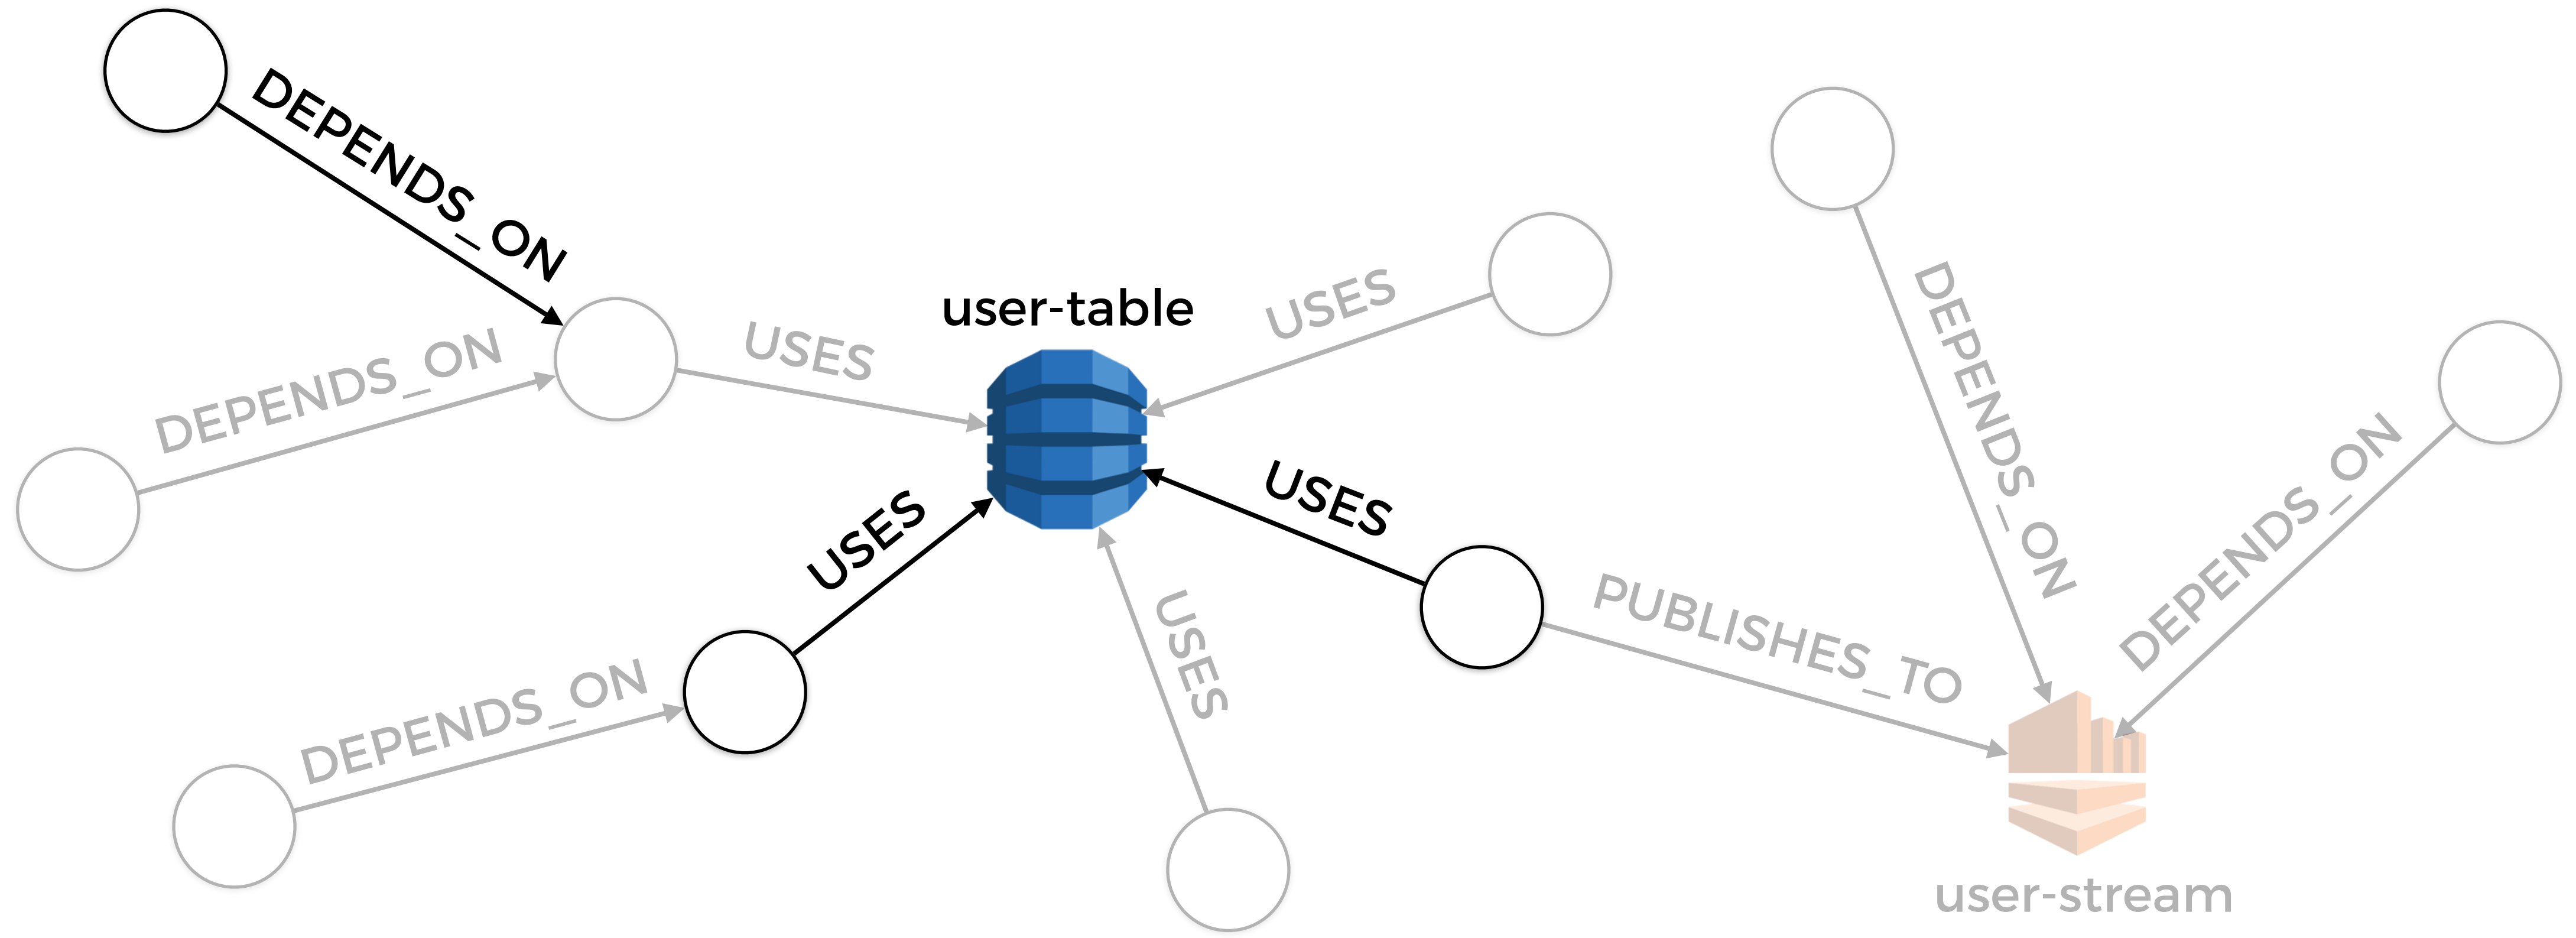




“what else is impacted by the throttled exceptions on user-table?”



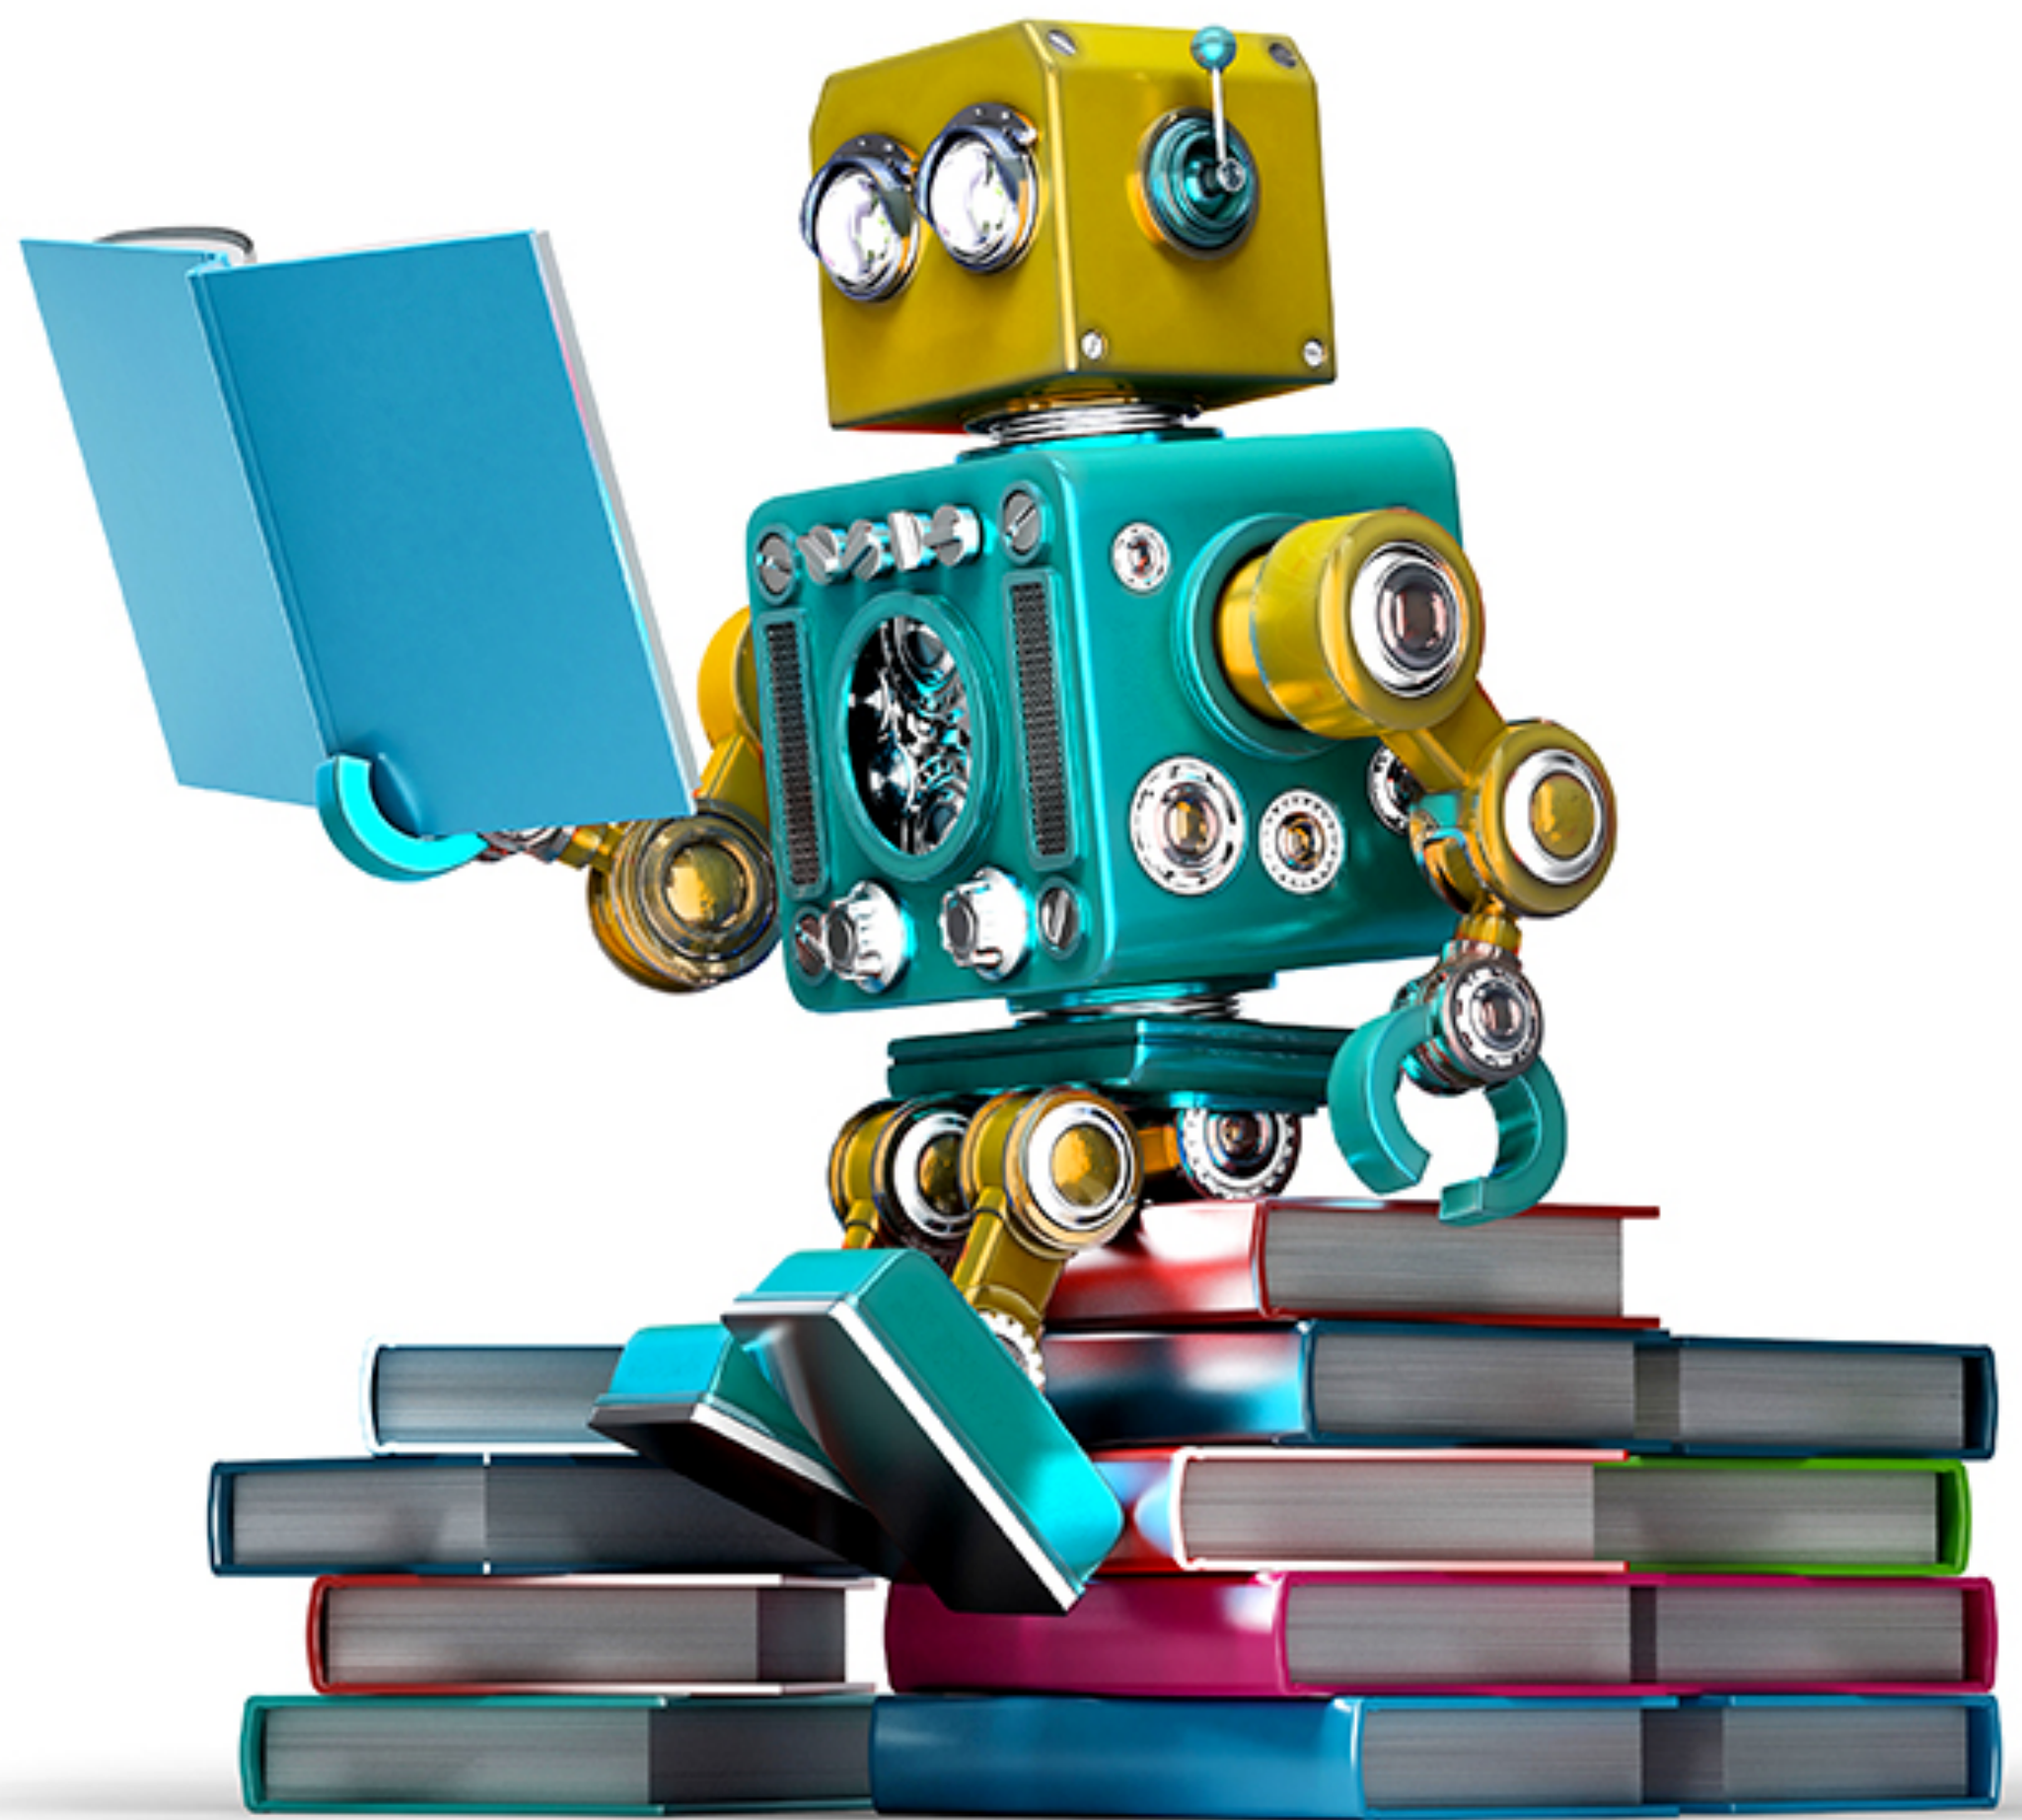
“what else is impacted by the throttled exceptions on user-table?”



wouldn't that be nice?



MACHINE LEARNING



use ML to auto-detect **erroneous** or **suspicious** behaviours, or to suggest possible **improvements**





Function [X] just performed an **unexpected write** against DynamoDB table [Y].

Should I...

[ignore it from now on](#)
[shut it down!!](#)



COMPOSE

Inbox (805)

Starred

Important

Sent Mail

Drafts (10)

Categories

Deleted Messages

Job Agencies

Notes

Personal

Receipts

Sent Messages

Travel

Twitter

Work

[production] Suspicious behaviour detected!



Inbox x



Observability Bot <bot@bestobservability.com>

1:18 AM (0 minutes ago)



to me

Hello Yan,

The function [X] has performed an **unexpected write** against DynamoDB table [Y], it has never done that before, is this intended?

Should I...

[ignore it from now on](#)
[shut it down!!!](#)



Click here to [Reply](#) or [Forward](#)

COMPOSE

Inbox (805)

Starred

Important

Sent Mail

Drafts (10)

Categories

Deleted Messages

Job Agencies

Notes

Personal

Receipts

Sent Messages

Travel

Twitter

Work

[production] Suspicious behaviour detected!



Inbox x



Observability Bot <bot@bestobservability.com>

1:18 AM (0 minutes ago)



to me

Hello Yan,

The function [X] has performed an **unexpected write** against DynamoDB table [Y], it has never done that before, is this intended?

Should I...

[ignore it from now on](#)
[shut it down!!!](#)

don't bother me about this again



Click here to [Reply](#) or [Forward](#)

COMPOSE

Inbox (805)

Starred

Important

Sent Mail

Drafts (10)

Categories

Deleted Messages

Job Agencies

Notes

Personal

Receipts

Sent Messages

Travel

Twitter

Work

[production] Suspicious behaviour detected!



Inbox x



Observability Bot <bot@bestobservability.com>

1:18 AM (0 minutes ago)



to me

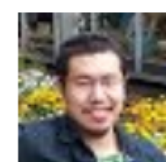
Hello Yan,

The function [X] has performed an **unexpected write** against DynamoDB table [Y], it has never done that before, is this intended?

Should I...

[ignore it from now on](#)
[shut it down!!!](#)

auto-modify IAM role with DENY rule

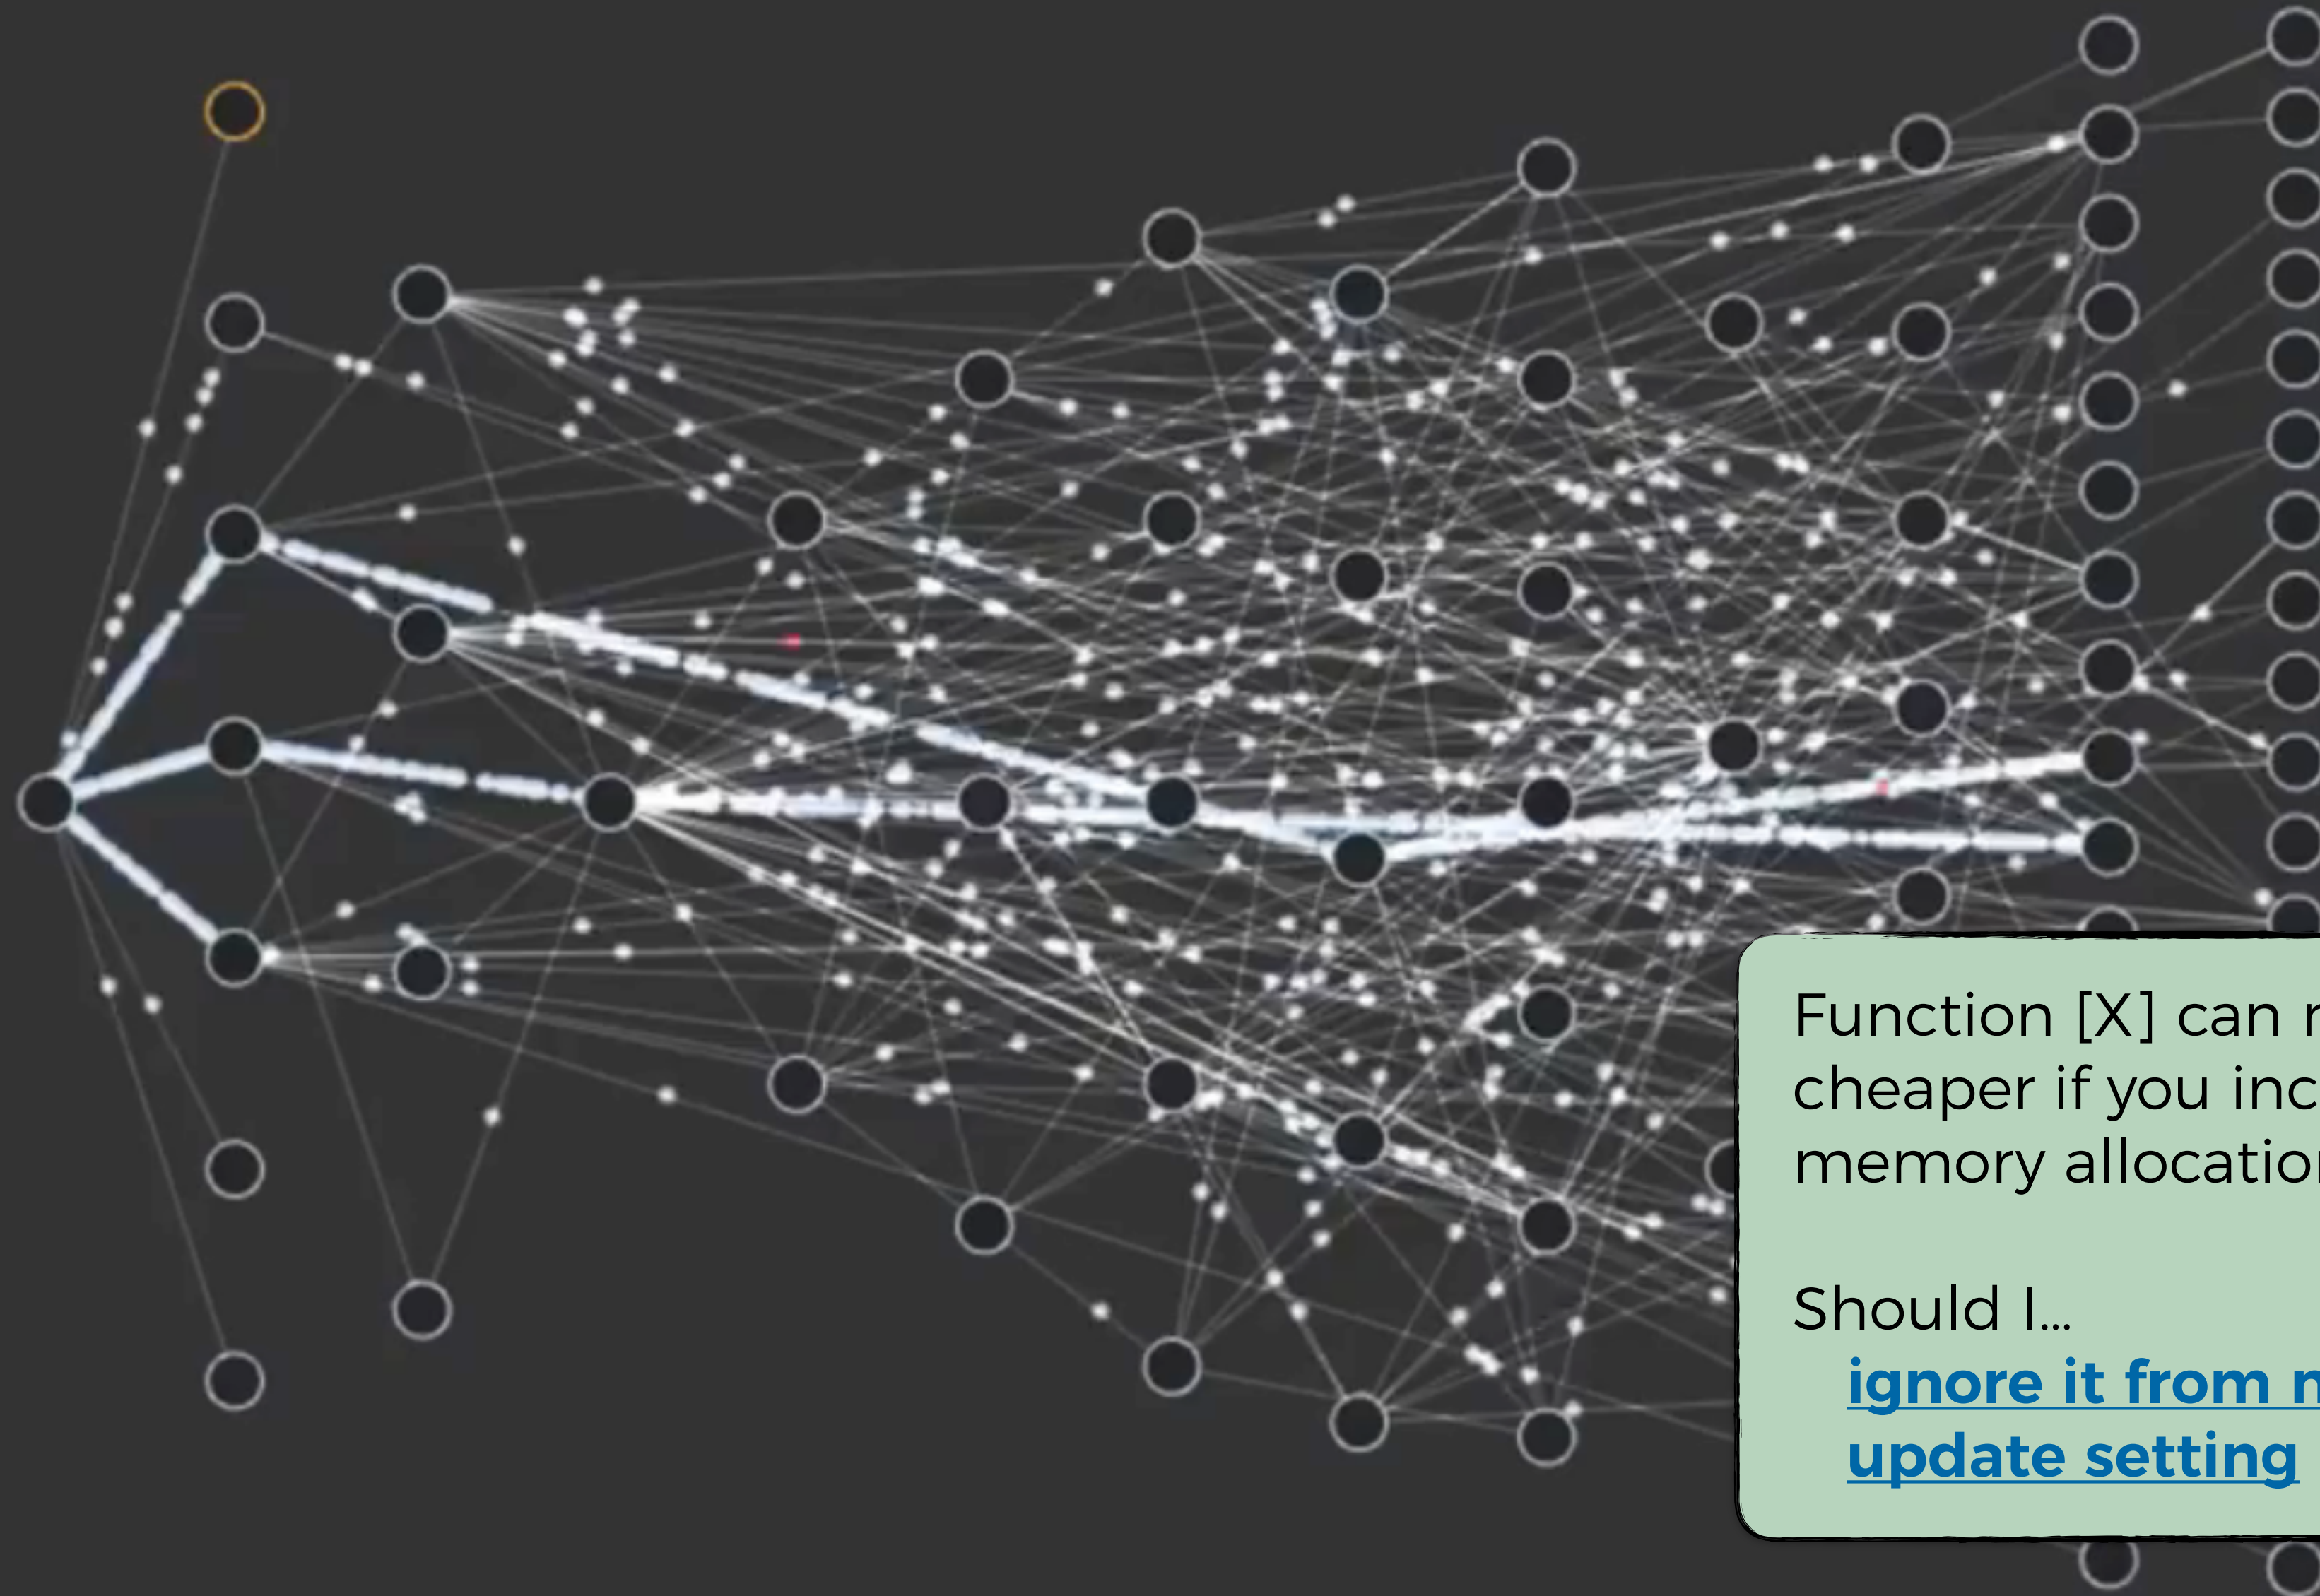


Click here to [Reply](#) or [Forward](#)



Function [X]'s performance has degraded since yesterday - 99% latency has gone up by 47% from 100ms to 147ms.





Function [X] can run faster & cheaper if you increase its memory allocation.

Should I...

[ignore it from now on](#)
[update setting](#)



A long-haired tabby cat is curled up and sleeping on a white, quilted mat. The mat is placed on a green, tufted sofa. A thought bubble originates from the cat's head, containing the text: "zzz... the future of.. zzz ... serverless observability.. zzz".

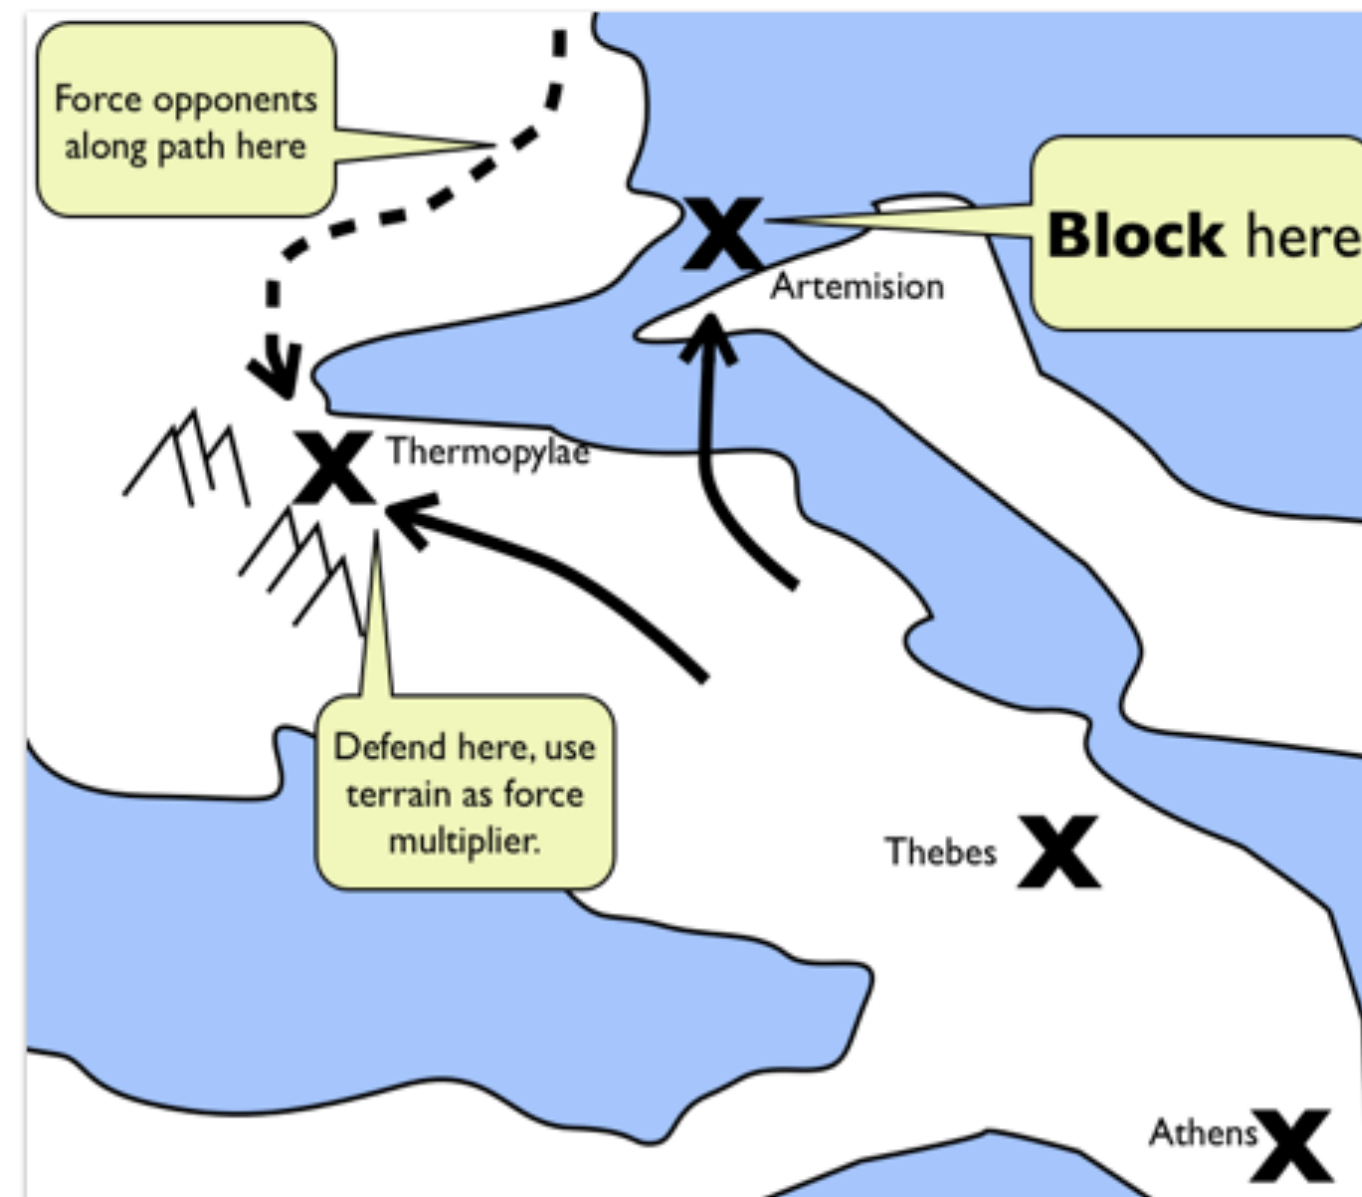
zzz... the future of.. zzz ...
serverless observability.. zzz





Simon Wardley

Which is MORE useful in determining strategy?



OR

<p>Strengths</p> <ul style="list-style-type: none"> A well trained Spartan army A high level of motivation not to become a Persian slave Most of the Persian army are mercenaries and slaves 	<p>Weaknesses</p> <ul style="list-style-type: none"> The rest of the Greeks aren't well trained The Ephors might stop the Spartans turning up A truck load of Persians ARE turning up
<p>Opportunities</p> <ul style="list-style-type: none"> Get rid of the Persians Get rid of the Spartans Become a Legend 	<p>Threats</p> <ul style="list-style-type: none"> Persians get rid of us The Oracle says a really dodgy film might be produced over 2,000 years later



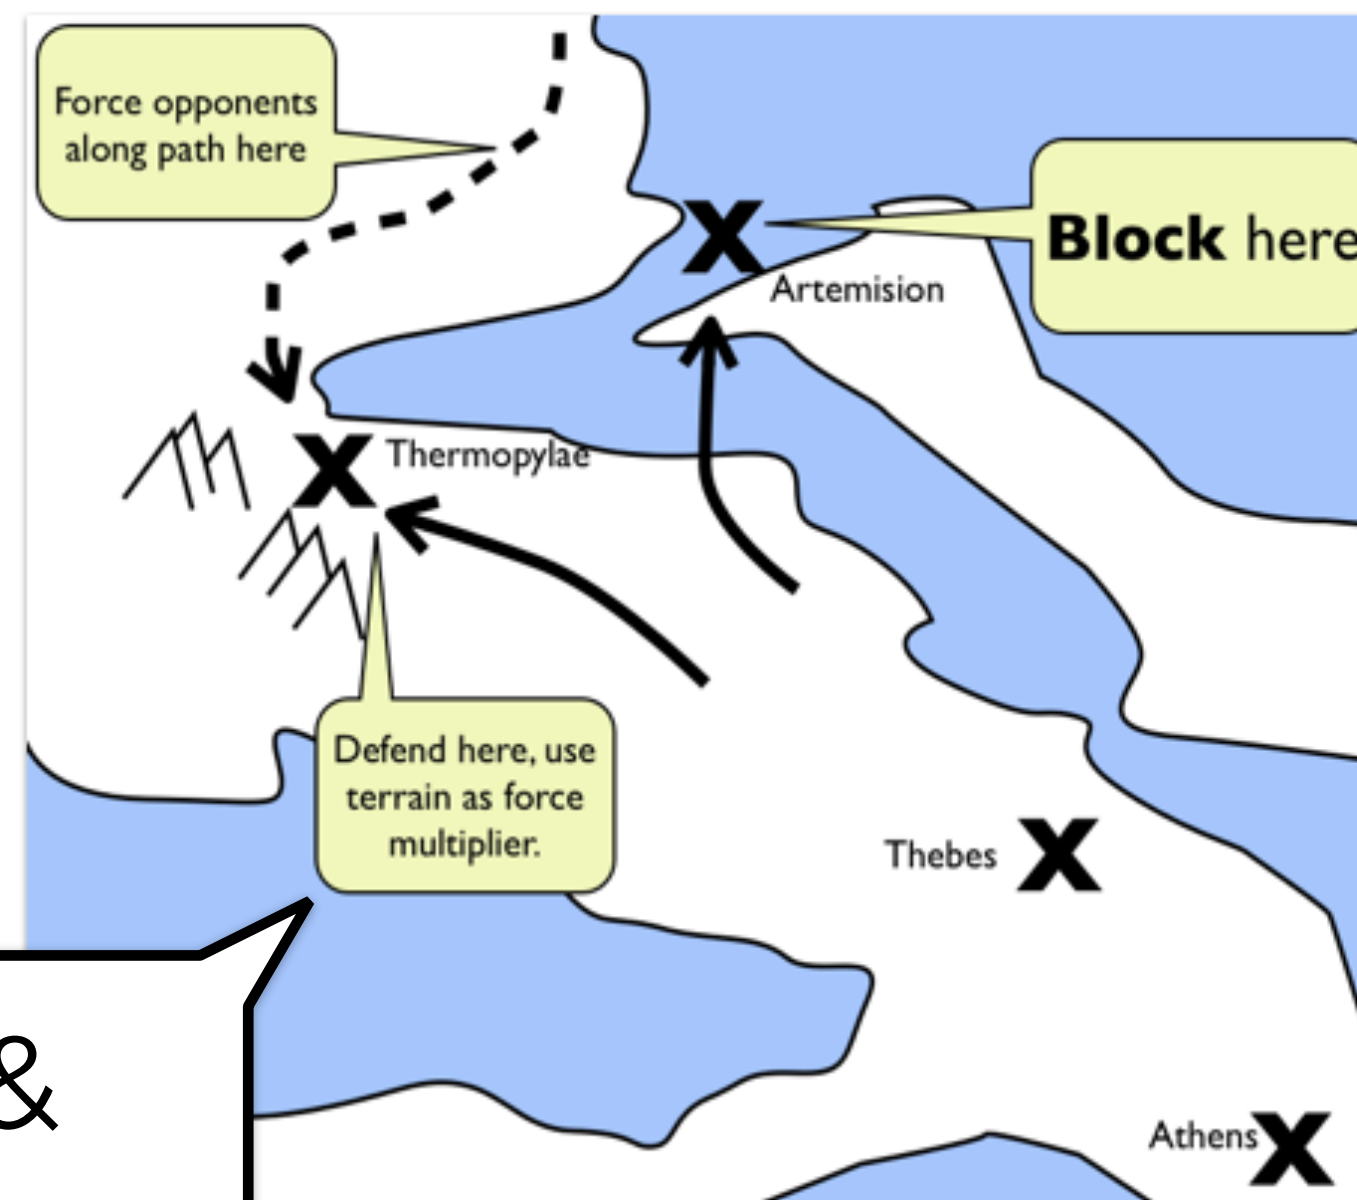
Which one do we use in business?

Which is MORE useful in determining strategy?



Simon Wardley

context & movement



OR

Strengths <ul style="list-style-type: none">A well trained Spartan armyA high level of motivation not to become a Persian slaveMost of the Persian army are mercenaries and slaves	Weaknesses <ul style="list-style-type: none">The rest of the Greeks aren't well trainedThe Ephors might stop the Spartans turning upA truck load of Persians ARE turning up
Opportunities <ul style="list-style-type: none">Get rid of the PersiansGet rid of the SpartansBecome a Legend	Threats <ul style="list-style-type: none">Persians get rid of usThe Oracle says a really dodgy film might be produced over 2,000 years later



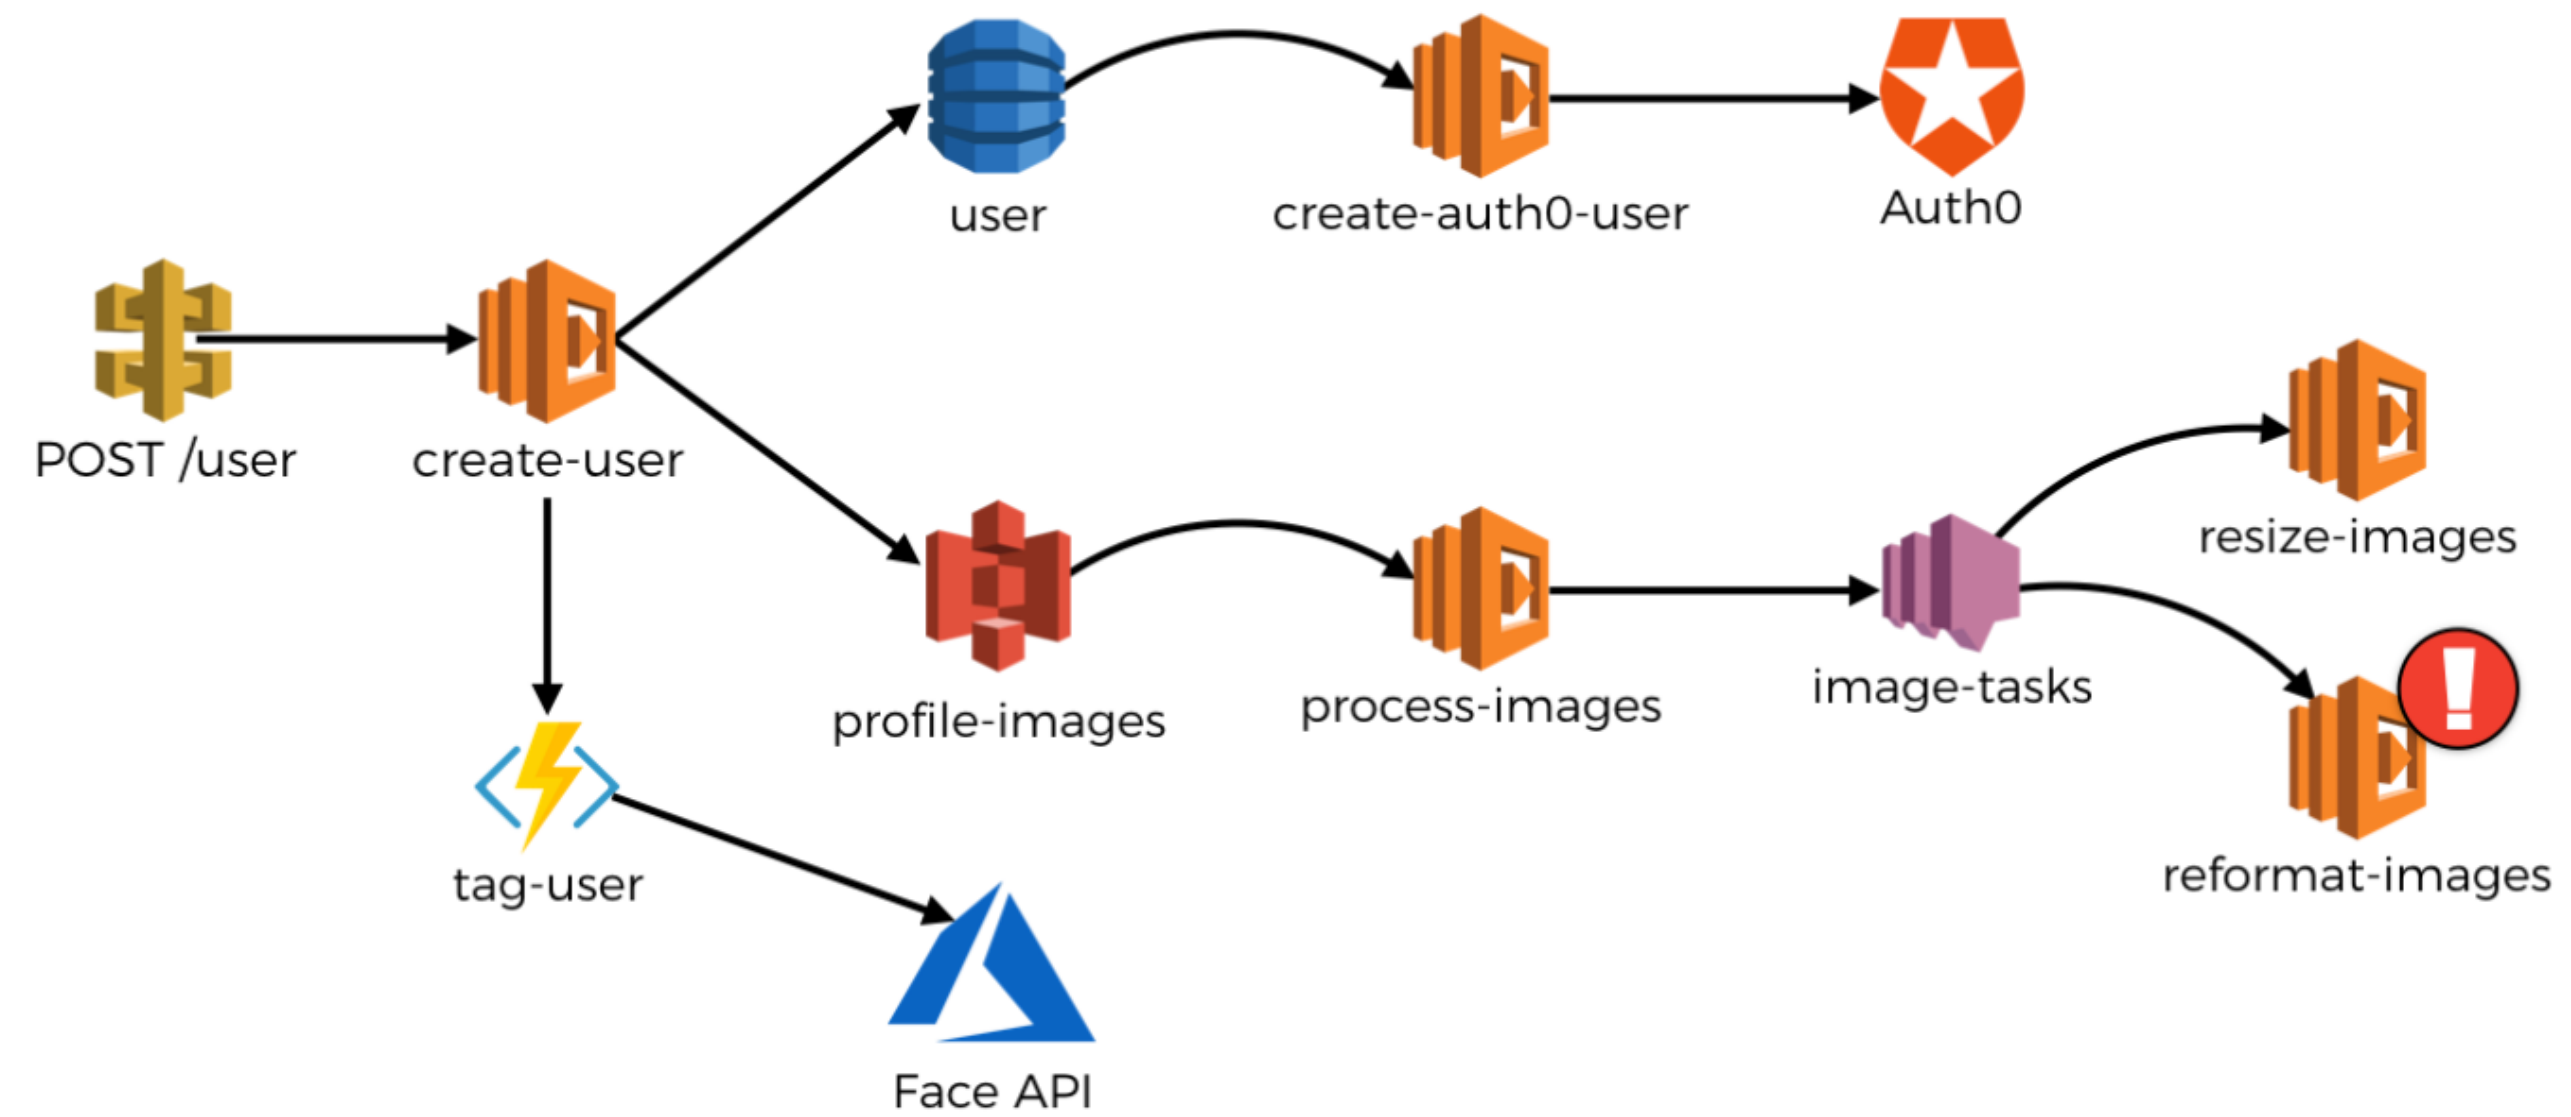
Which one do we use in business?



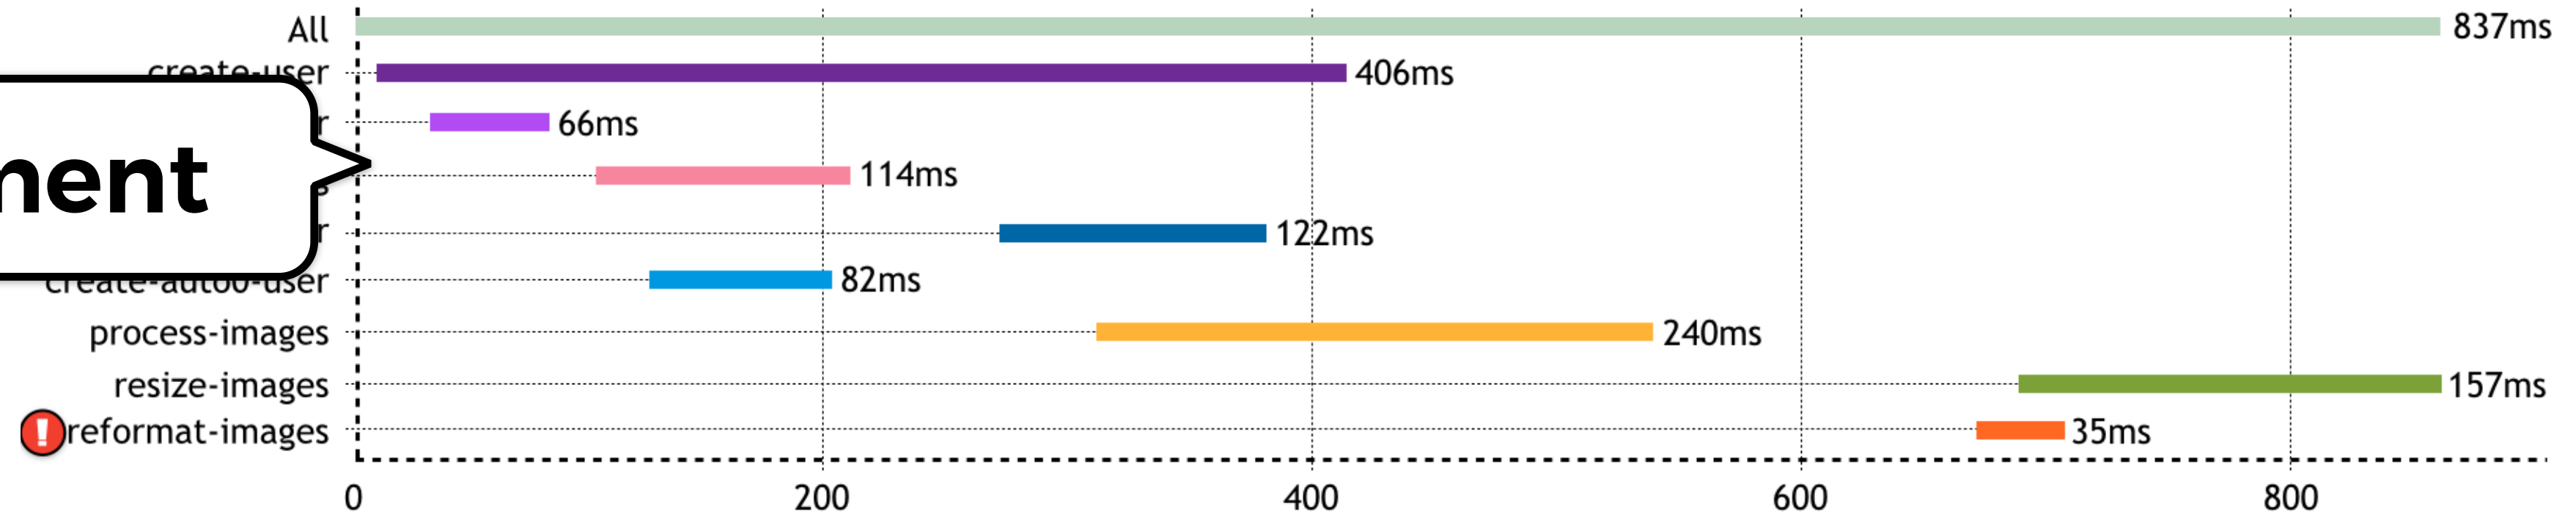
*However, I would argue that the health of the system no longer matters. We've entered an era where what matters is the health of each individual event, or each individual user's experience, or each shopping cart's experience (or other high cardinality dimensions). With distributed systems you don't care about the health of the system, you care about the **health of the event or the slice.** ””*

- Charity Majors <http://bit.ly/2E2QngU>

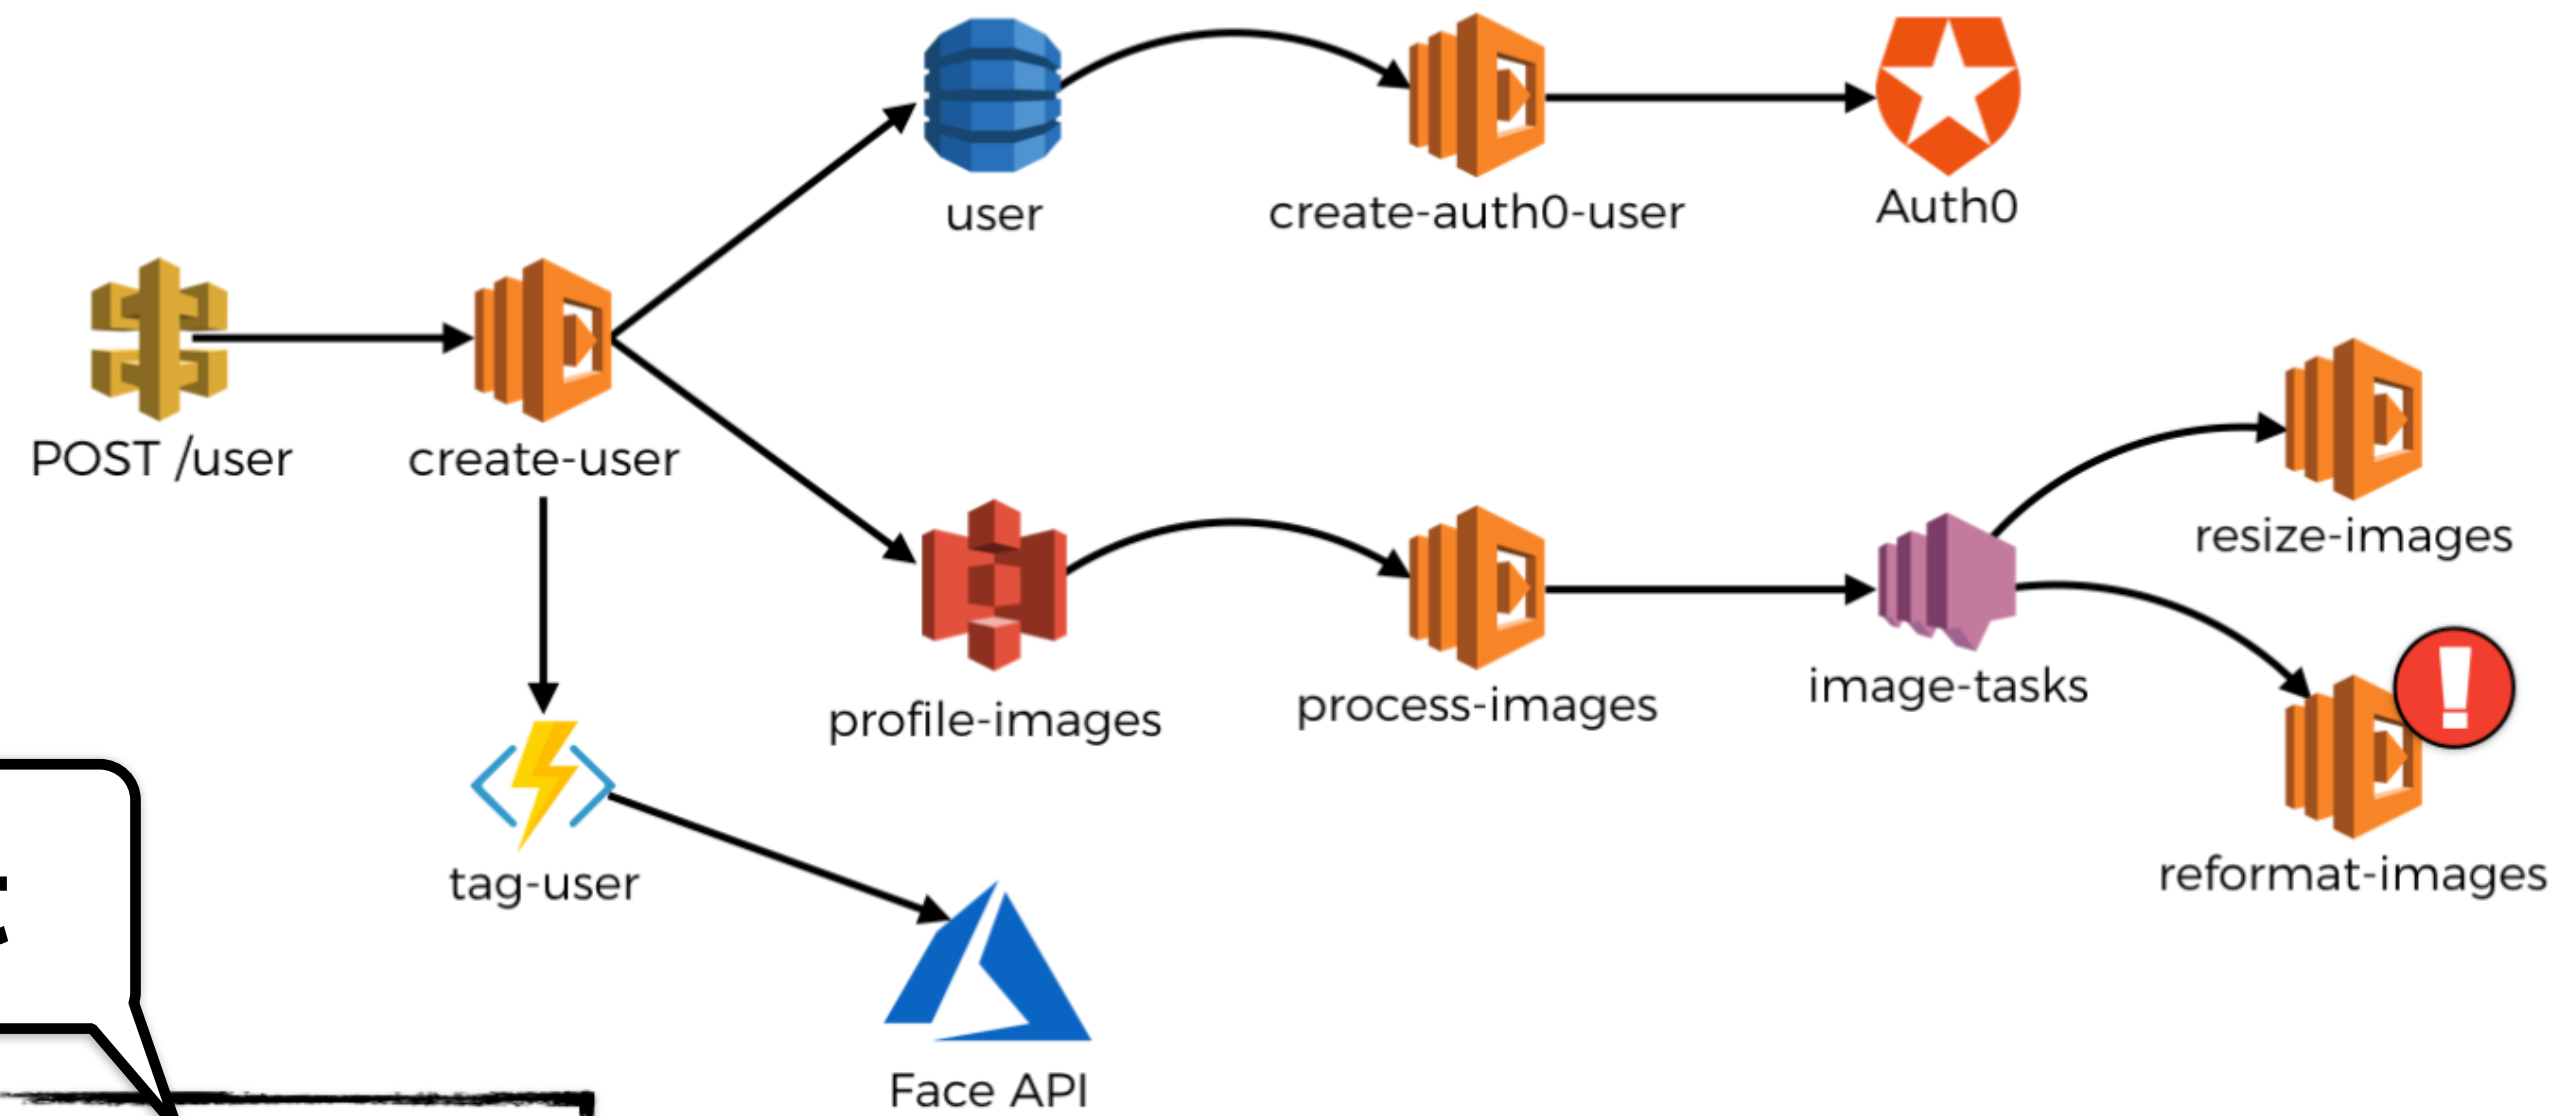
“one user action/vertical slice through the system”



movement



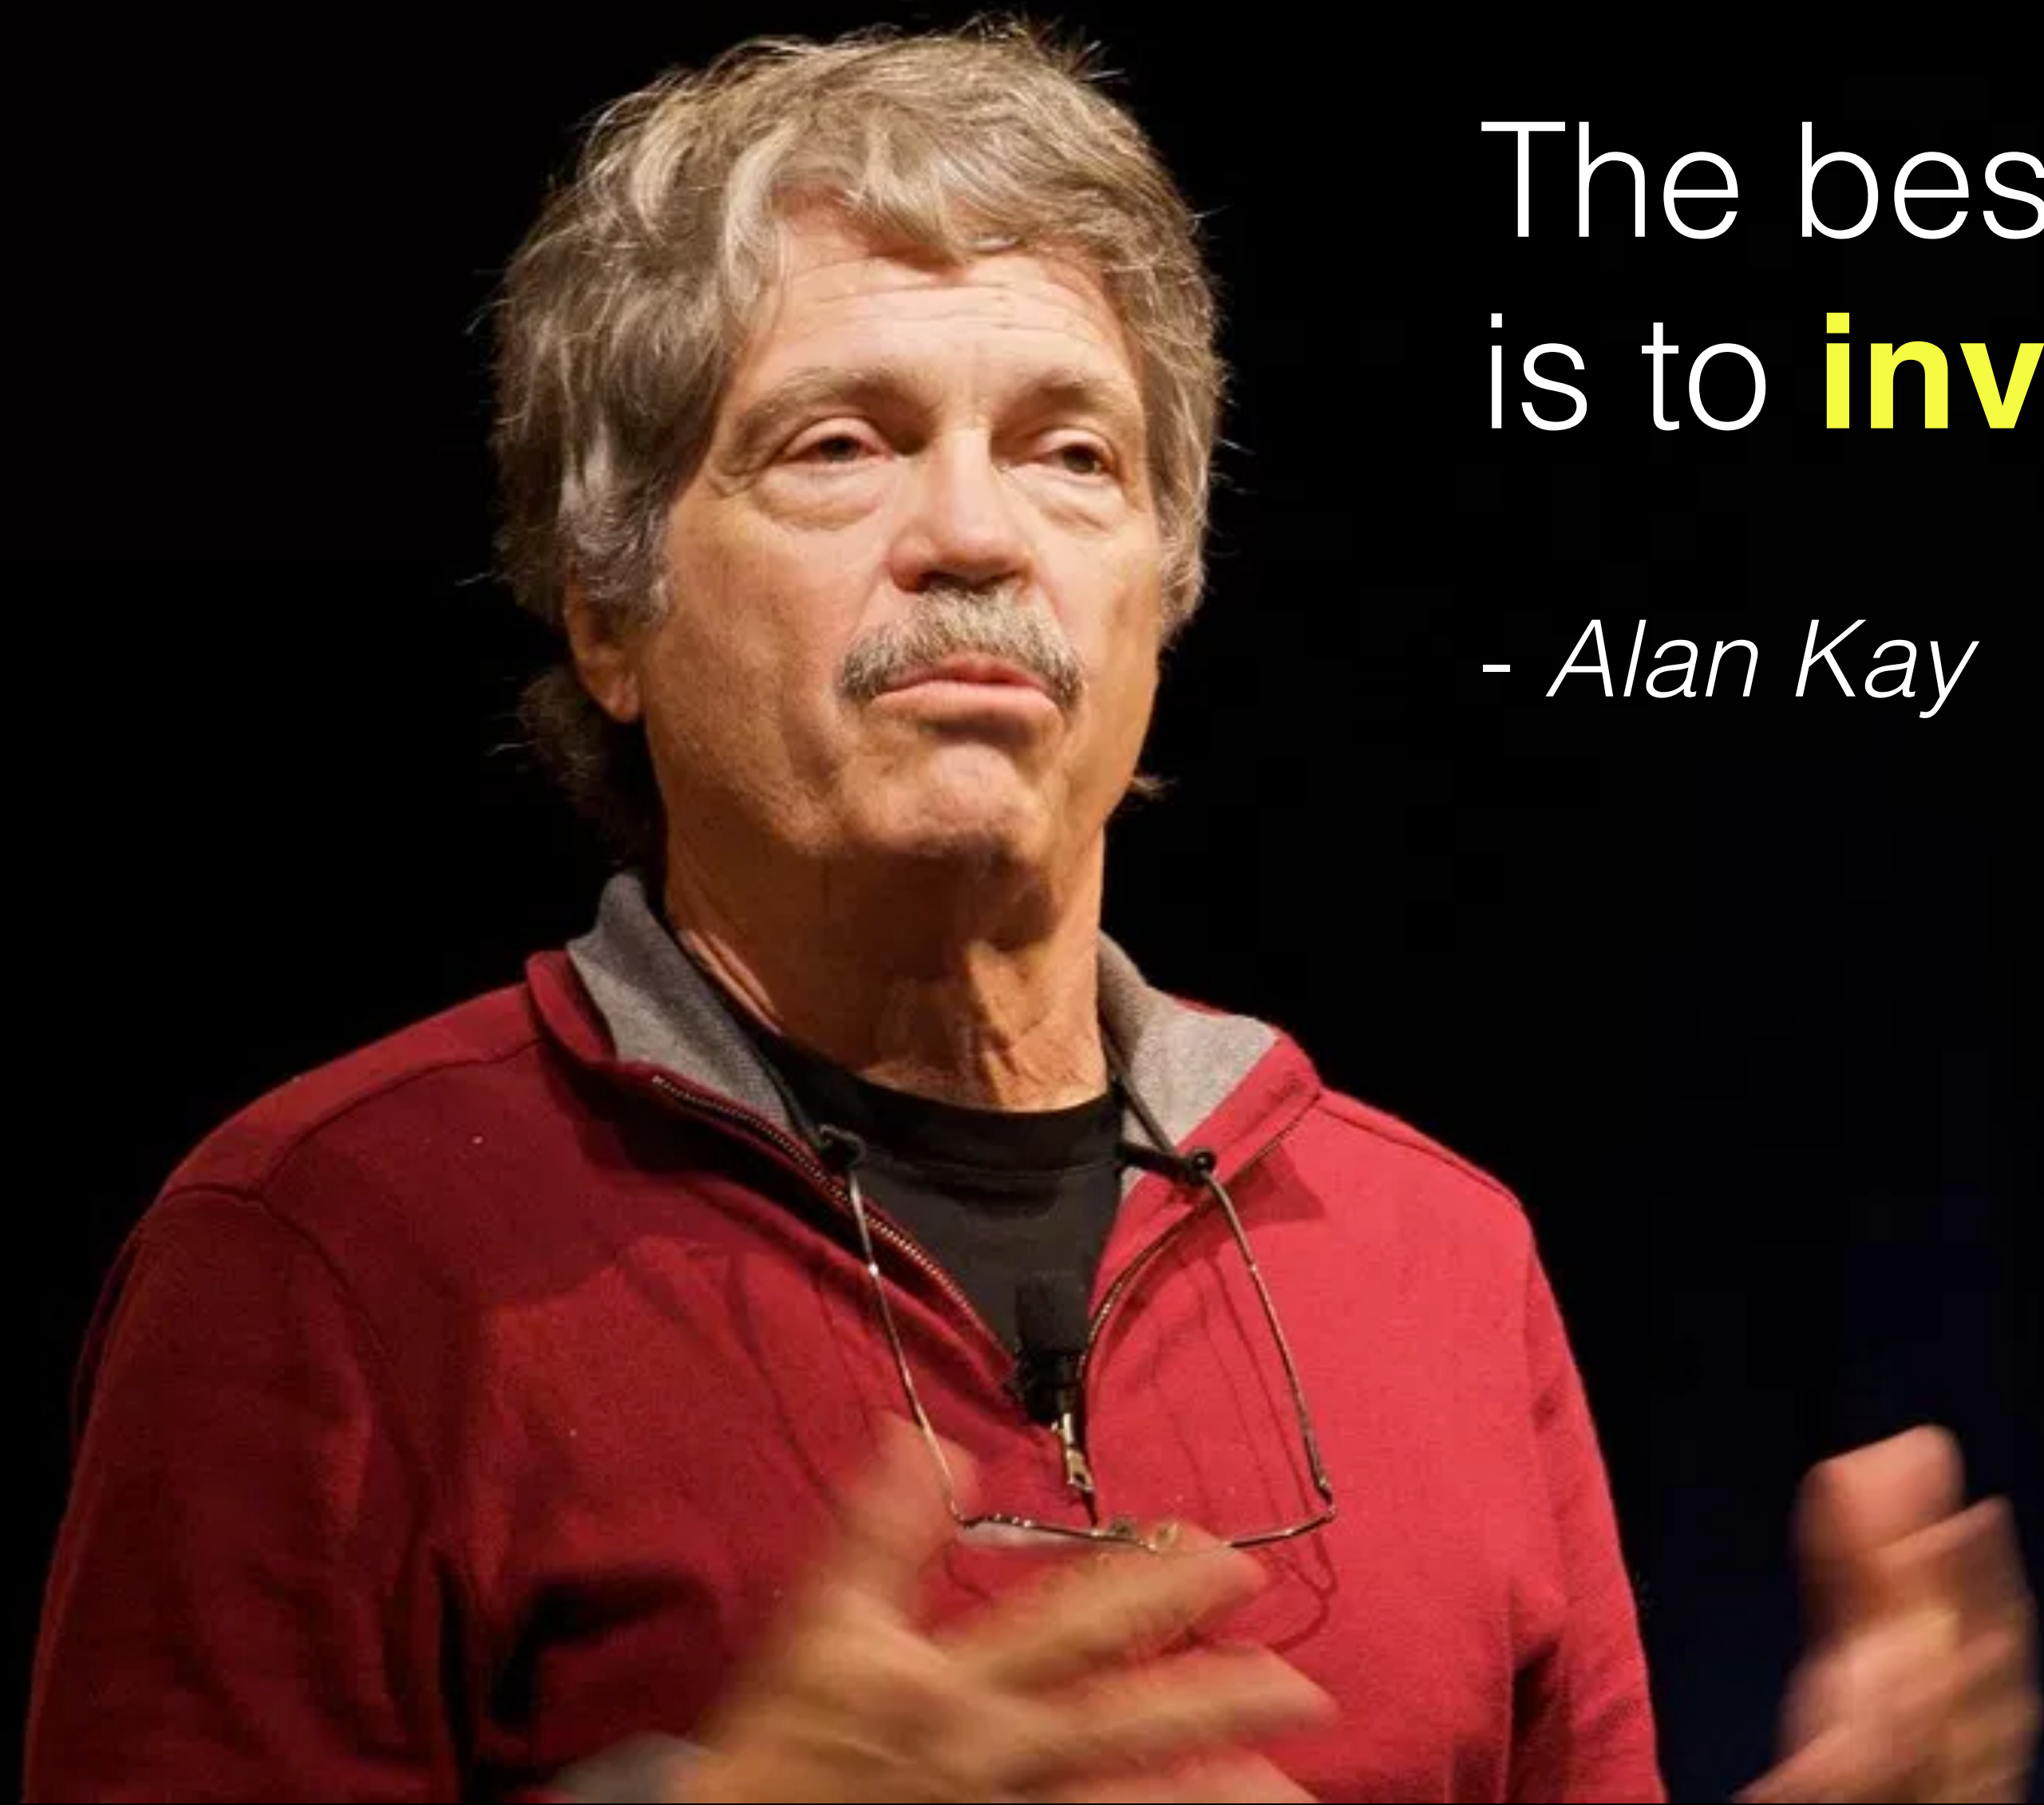
movement



context

Logs **Input/Output**

timestamp	component	level	message
2018/01/25 20:51:23.188	POST /user	debug	incoming request...
2018/01/25 20:51:23.201	create-user	debug	saving user [theburningmonk] in the [user] table...
2018/01/25 20:51:23.215	create-user	debug	saved user [theburningmonk] in the [user] table
2018/01/25 20:51:23.521	tag-user	debug	tagging user [theburningmonk] with Azure Face API...



The best way to predict the future
is to **invent it.**

- *Alan Kay*

**The best way to invent
the future is to inception
someone else to do it.**



- *me*





FULL OBSERVABILITY FOR AWS LAMBDA

Instrument and profile your functions with zero overhead.

Gain visibility to identify and resolve issues faster.

[Sign up for early access](#)



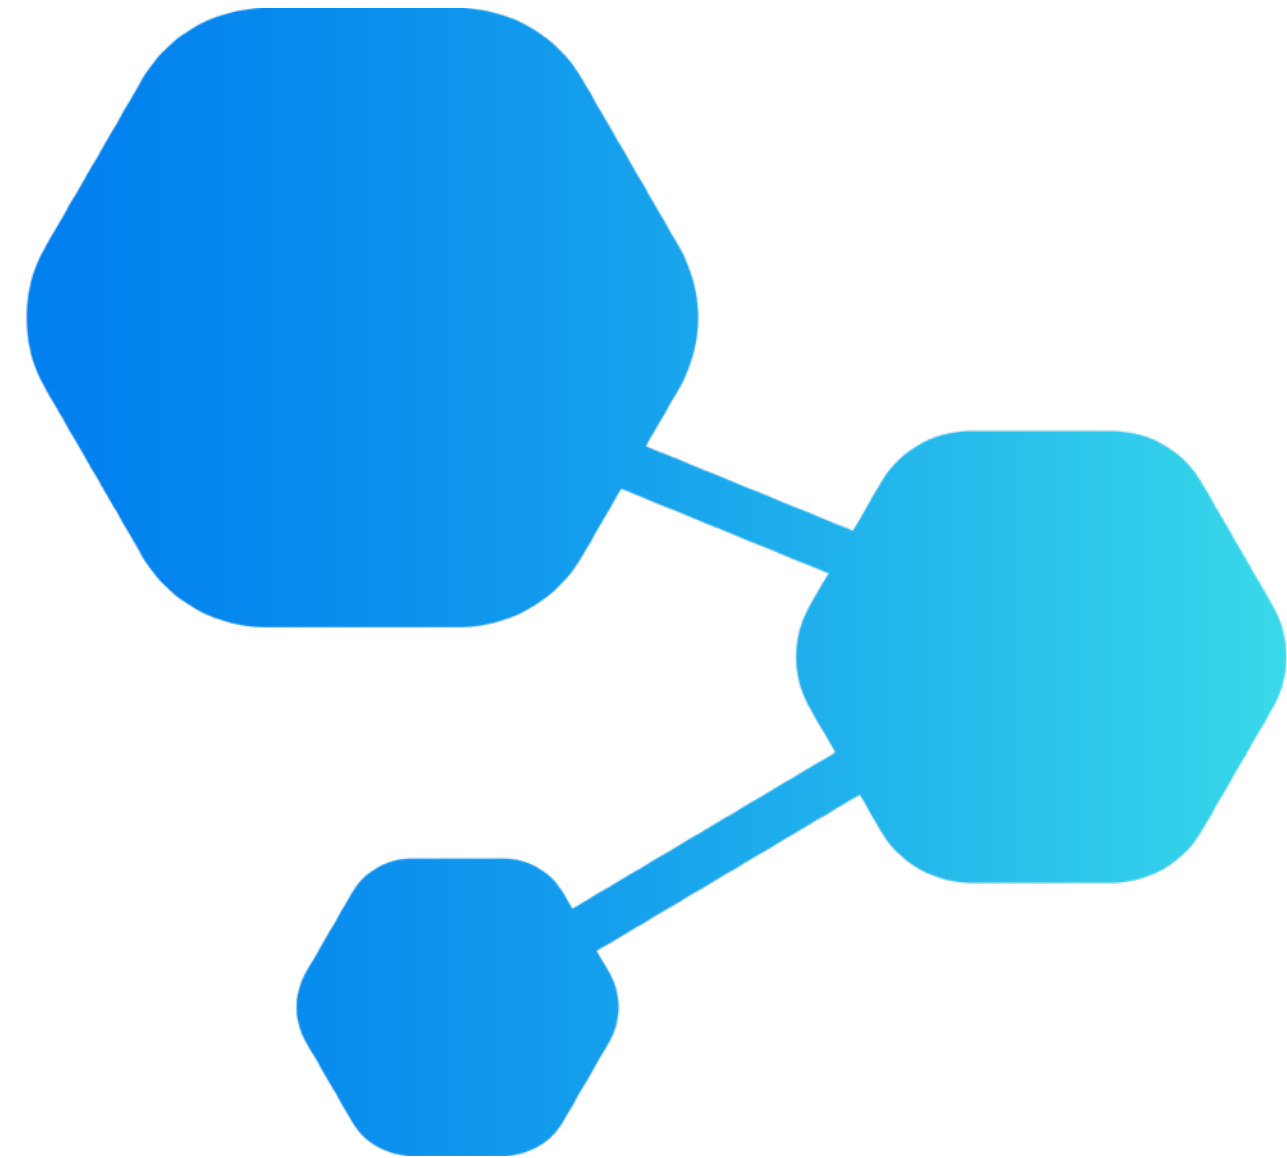
Serkan Özal
[@serkan_ozal](#)



Nitzan Shapira
@nitzanshapira



Ran Ribenzaft
@ranrib



Epsagon

SEE INSIDE YOUR LAMBDA FUNCTIONS



Adam Johnson
@adjohn



Erica Windisch
@ewindisch

Try it Free

Scroll Down to Find Out How





Charity Majors
@mipsytipsy



Cindy Sridharan
@copyconstruct



Liz Fong-Jones
@lizthegrey



JBD
@rakyll



Erica Windisch
@ewindisch



Production-Ready Serverless **MEAP**

Operational Best Practices

- API Gateway and Kinesis
- Authentication & authorisation (IAM, Cognito)
- Testing
- Running & Debugging functions locally
- Log aggregation
- Monitoring & Alerting
- X-Ray
- Correlation IDs
- CI/CD
- Performance and Cost optimisation
- Error Handling
- Configuration management
- VPC
- Security
- Leading practices (API Gateway, Kinesis, Lambda)
- Canary deployments

**get 40% off
with code:
ytcui**

<http://bit.ly/production-ready-serverless>