
AWS
re:Invent

ANT330

Intuit: Moving from monitoring to observability using Amazon ES

Balaji Arunachalam
Director of Engineering
Intuit Inc.

Nagaraja Tantry
Staff Software Engineer
Intuit Inc.

Ian Swanson
General Manager
Amazon Web Services

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws



intuit

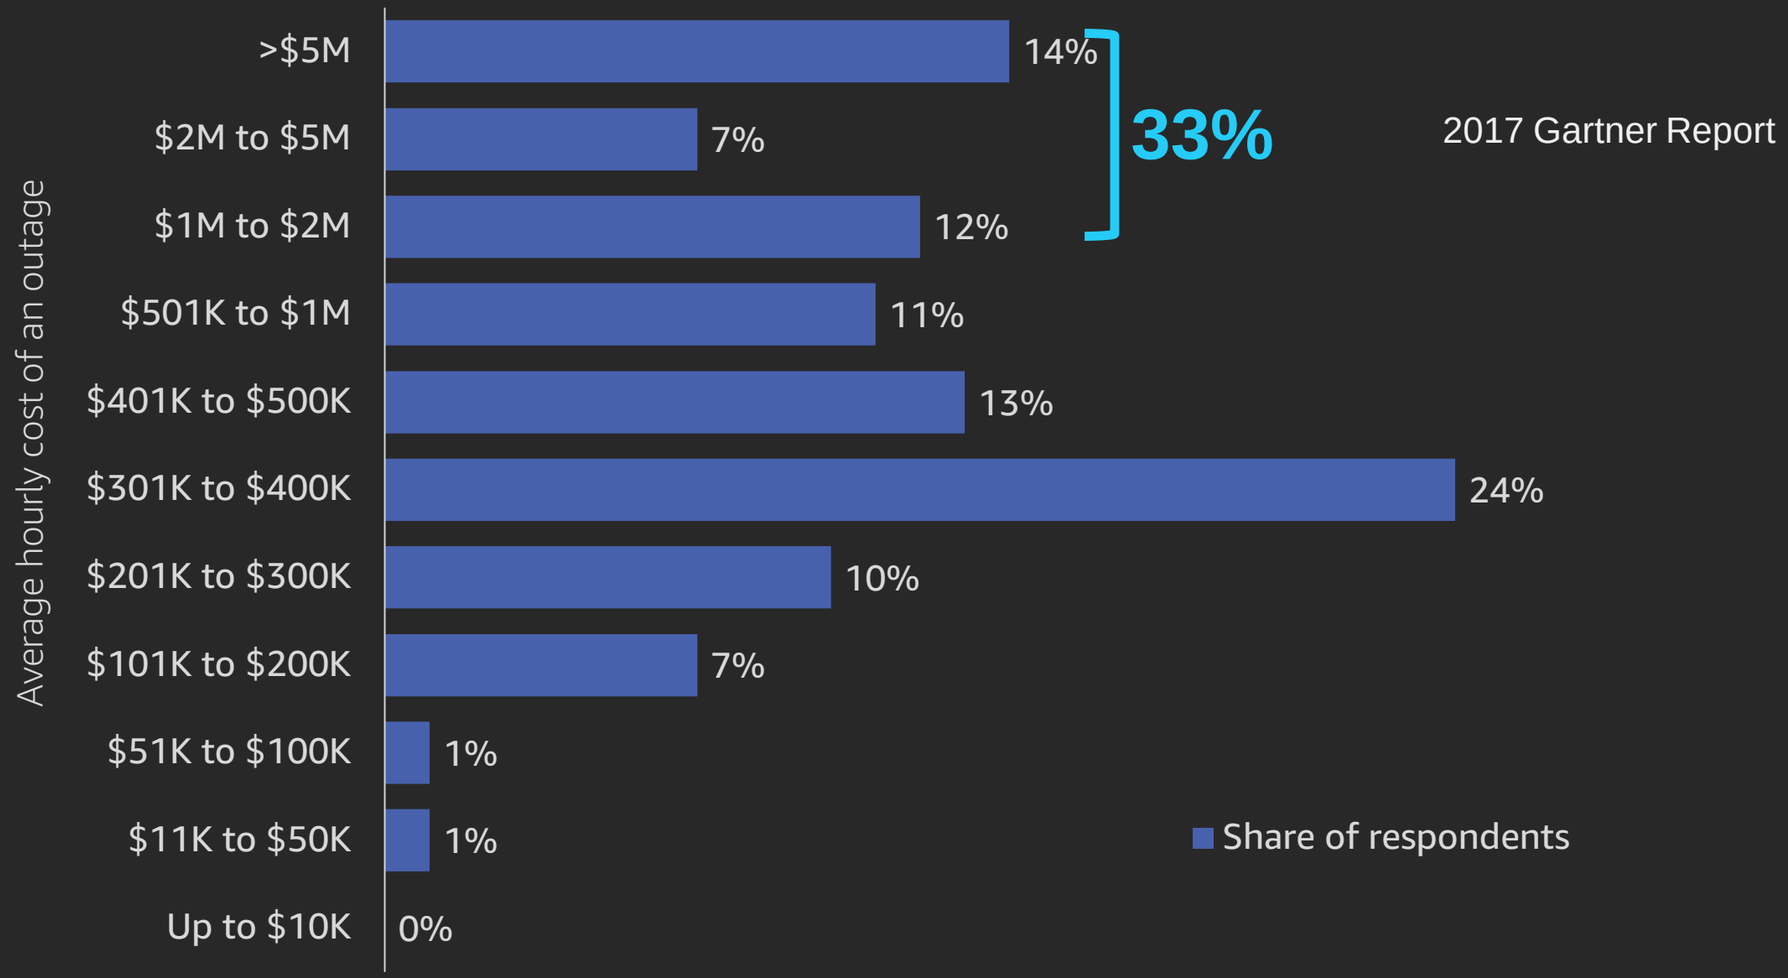
turbotax quickbooks mint

Who we serve

Consumers
Small businesses
Self-employed

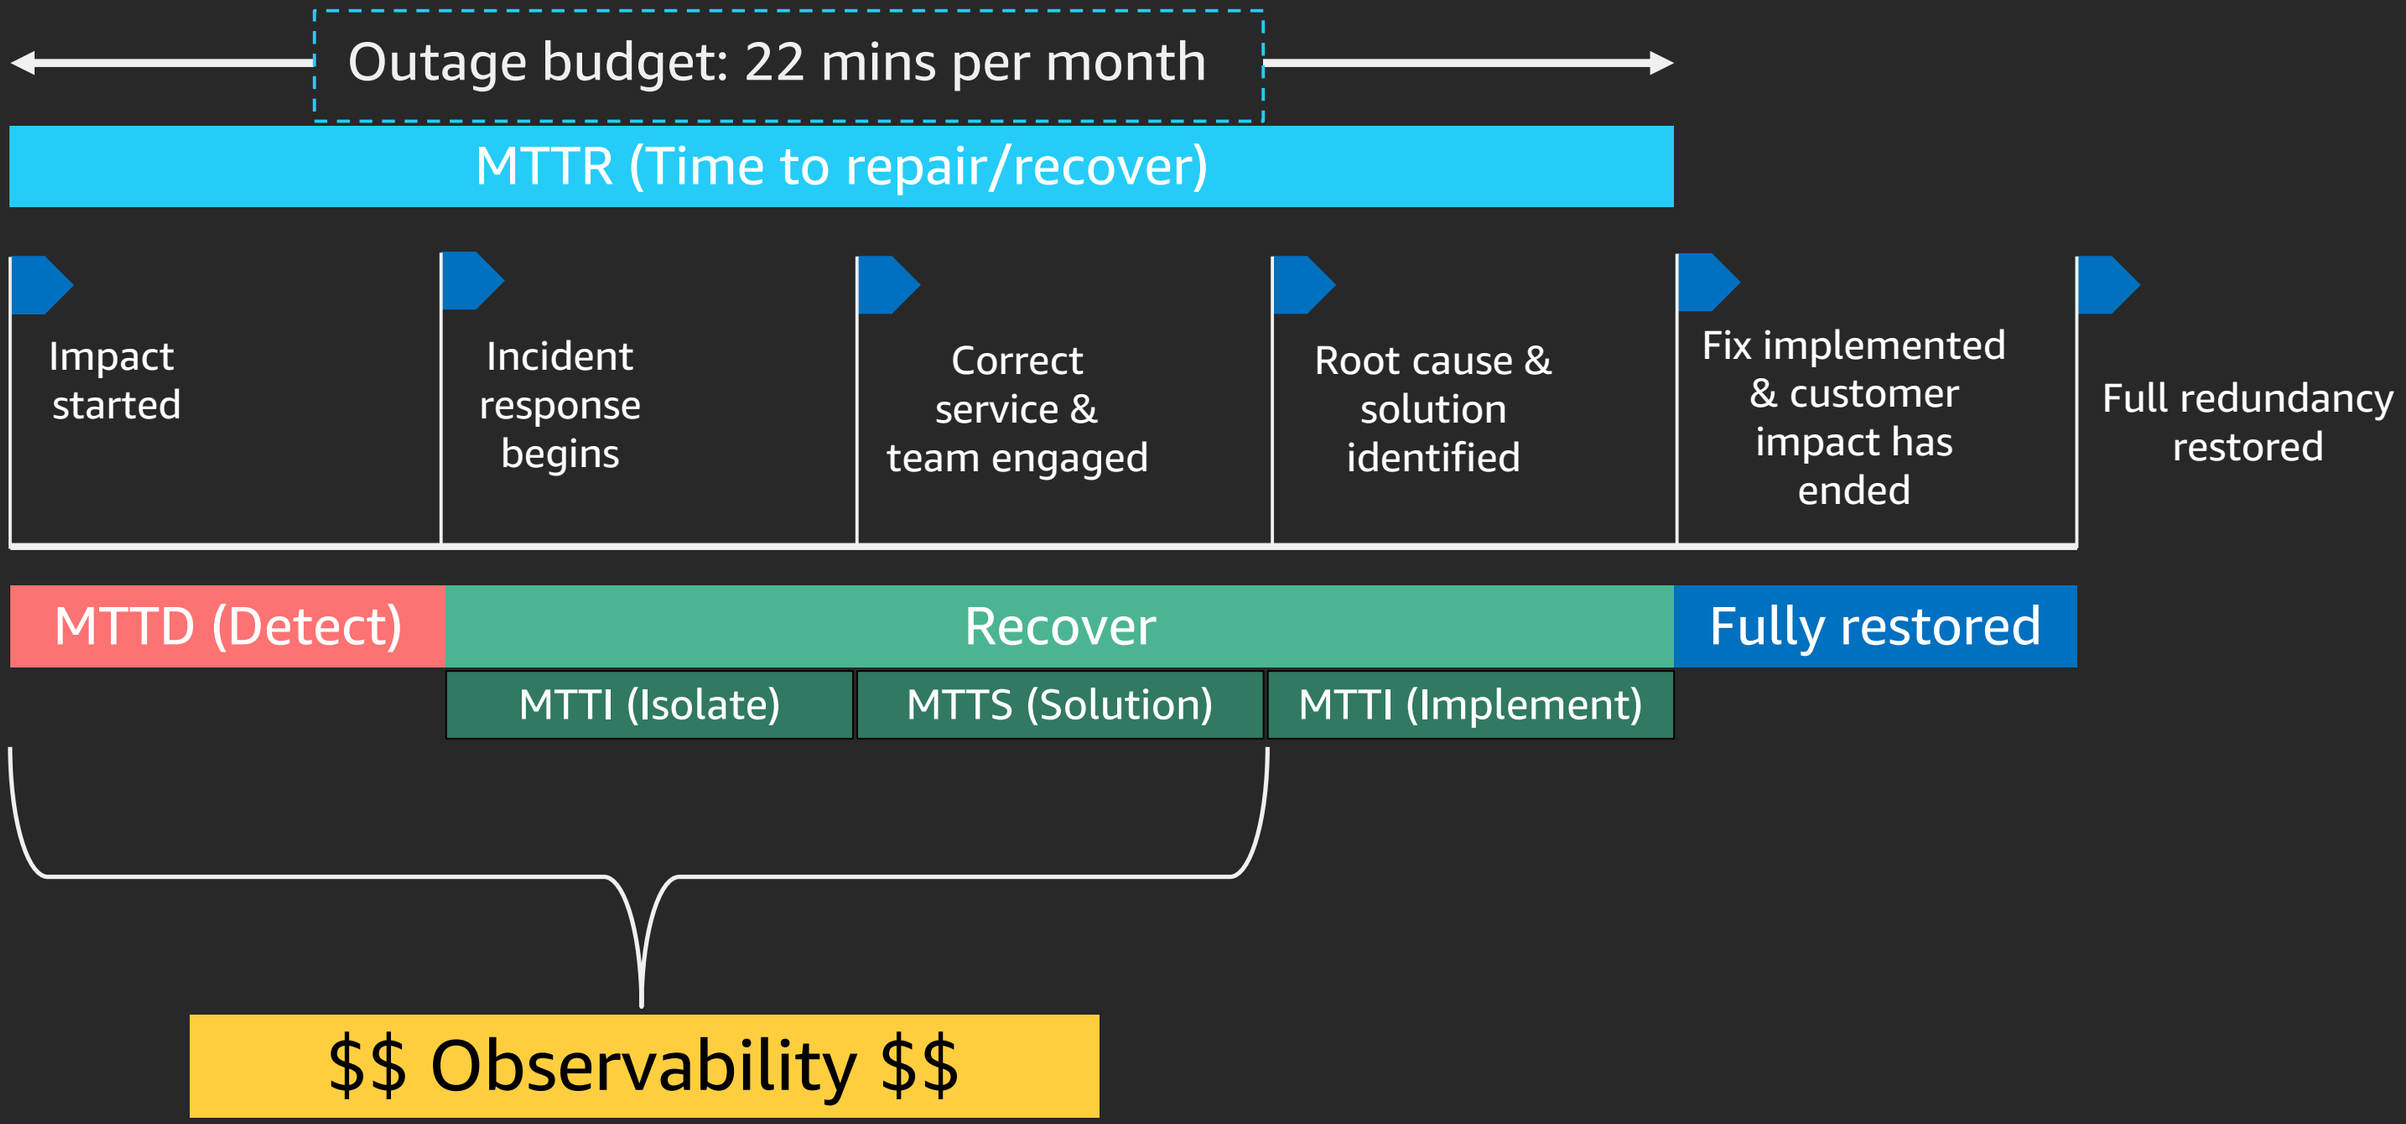
Hourly cost of an outage

33% of companies pay +$1M per hour of outage

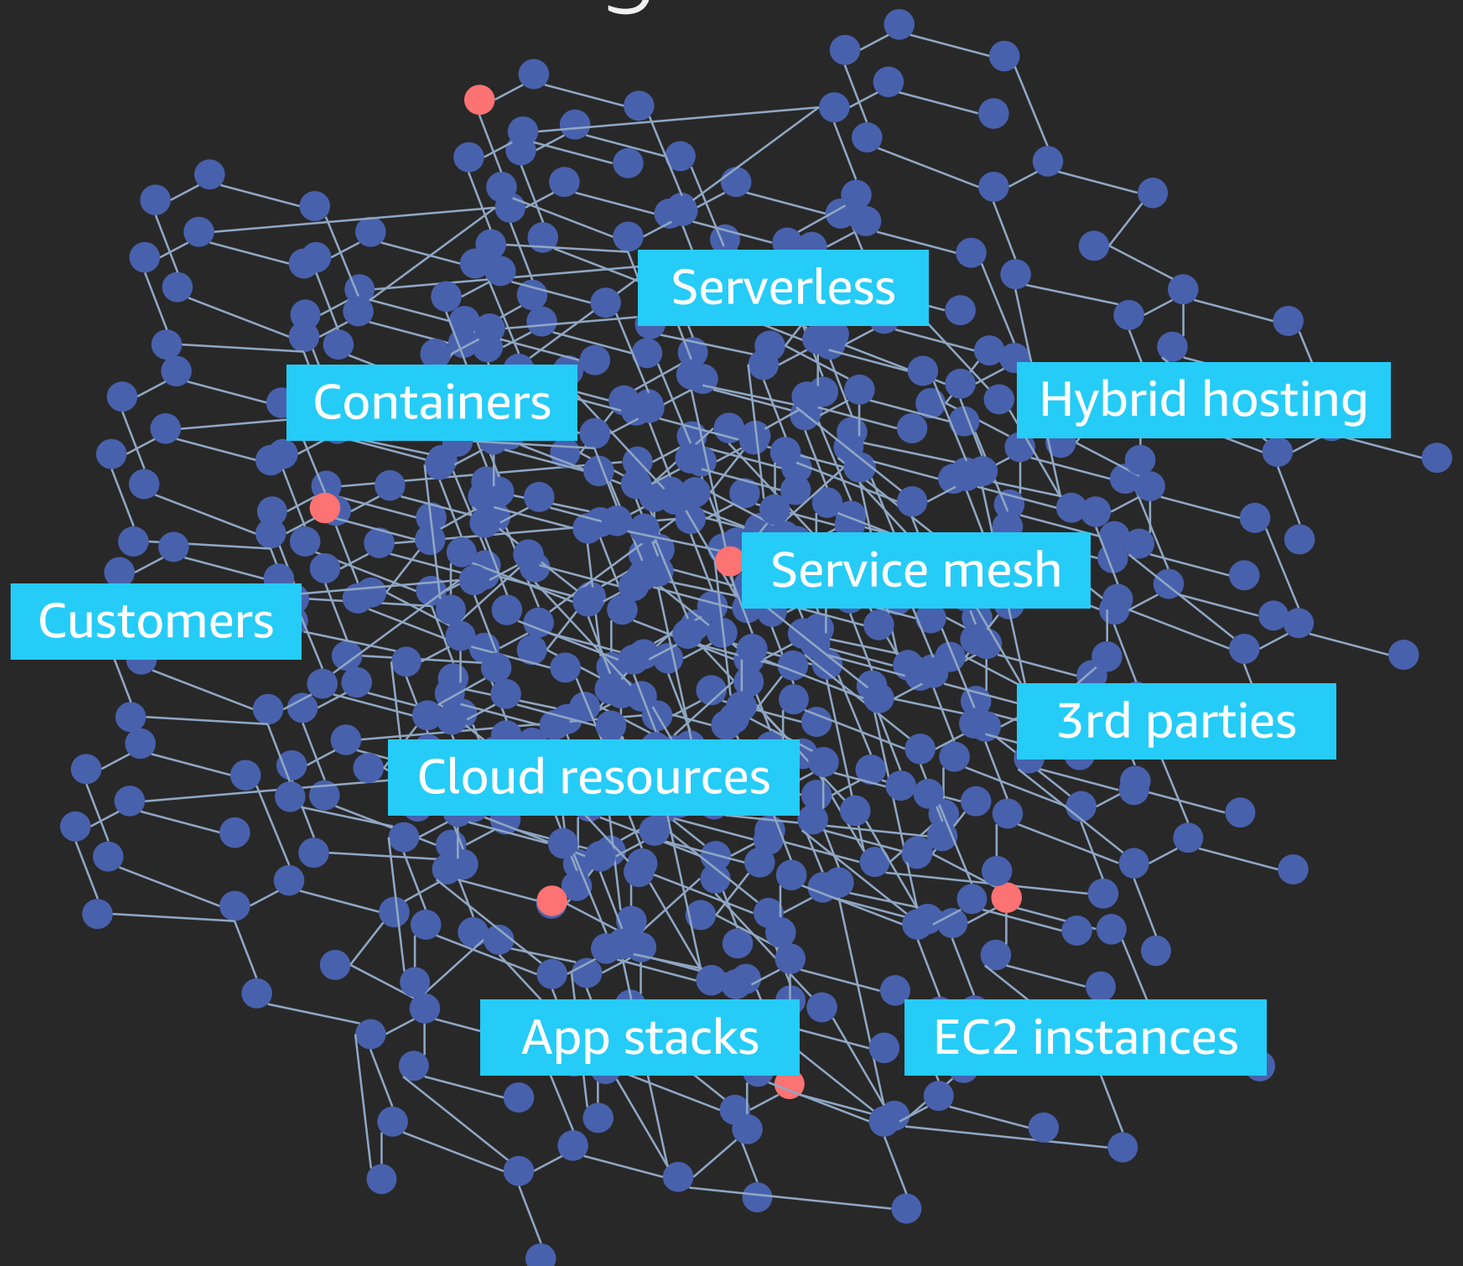


Source: Statista

Role of observability in reducing outage duration



Case of change at Intuit



Complex environment – 1.5 million PPS metrics (200% YoY growth)

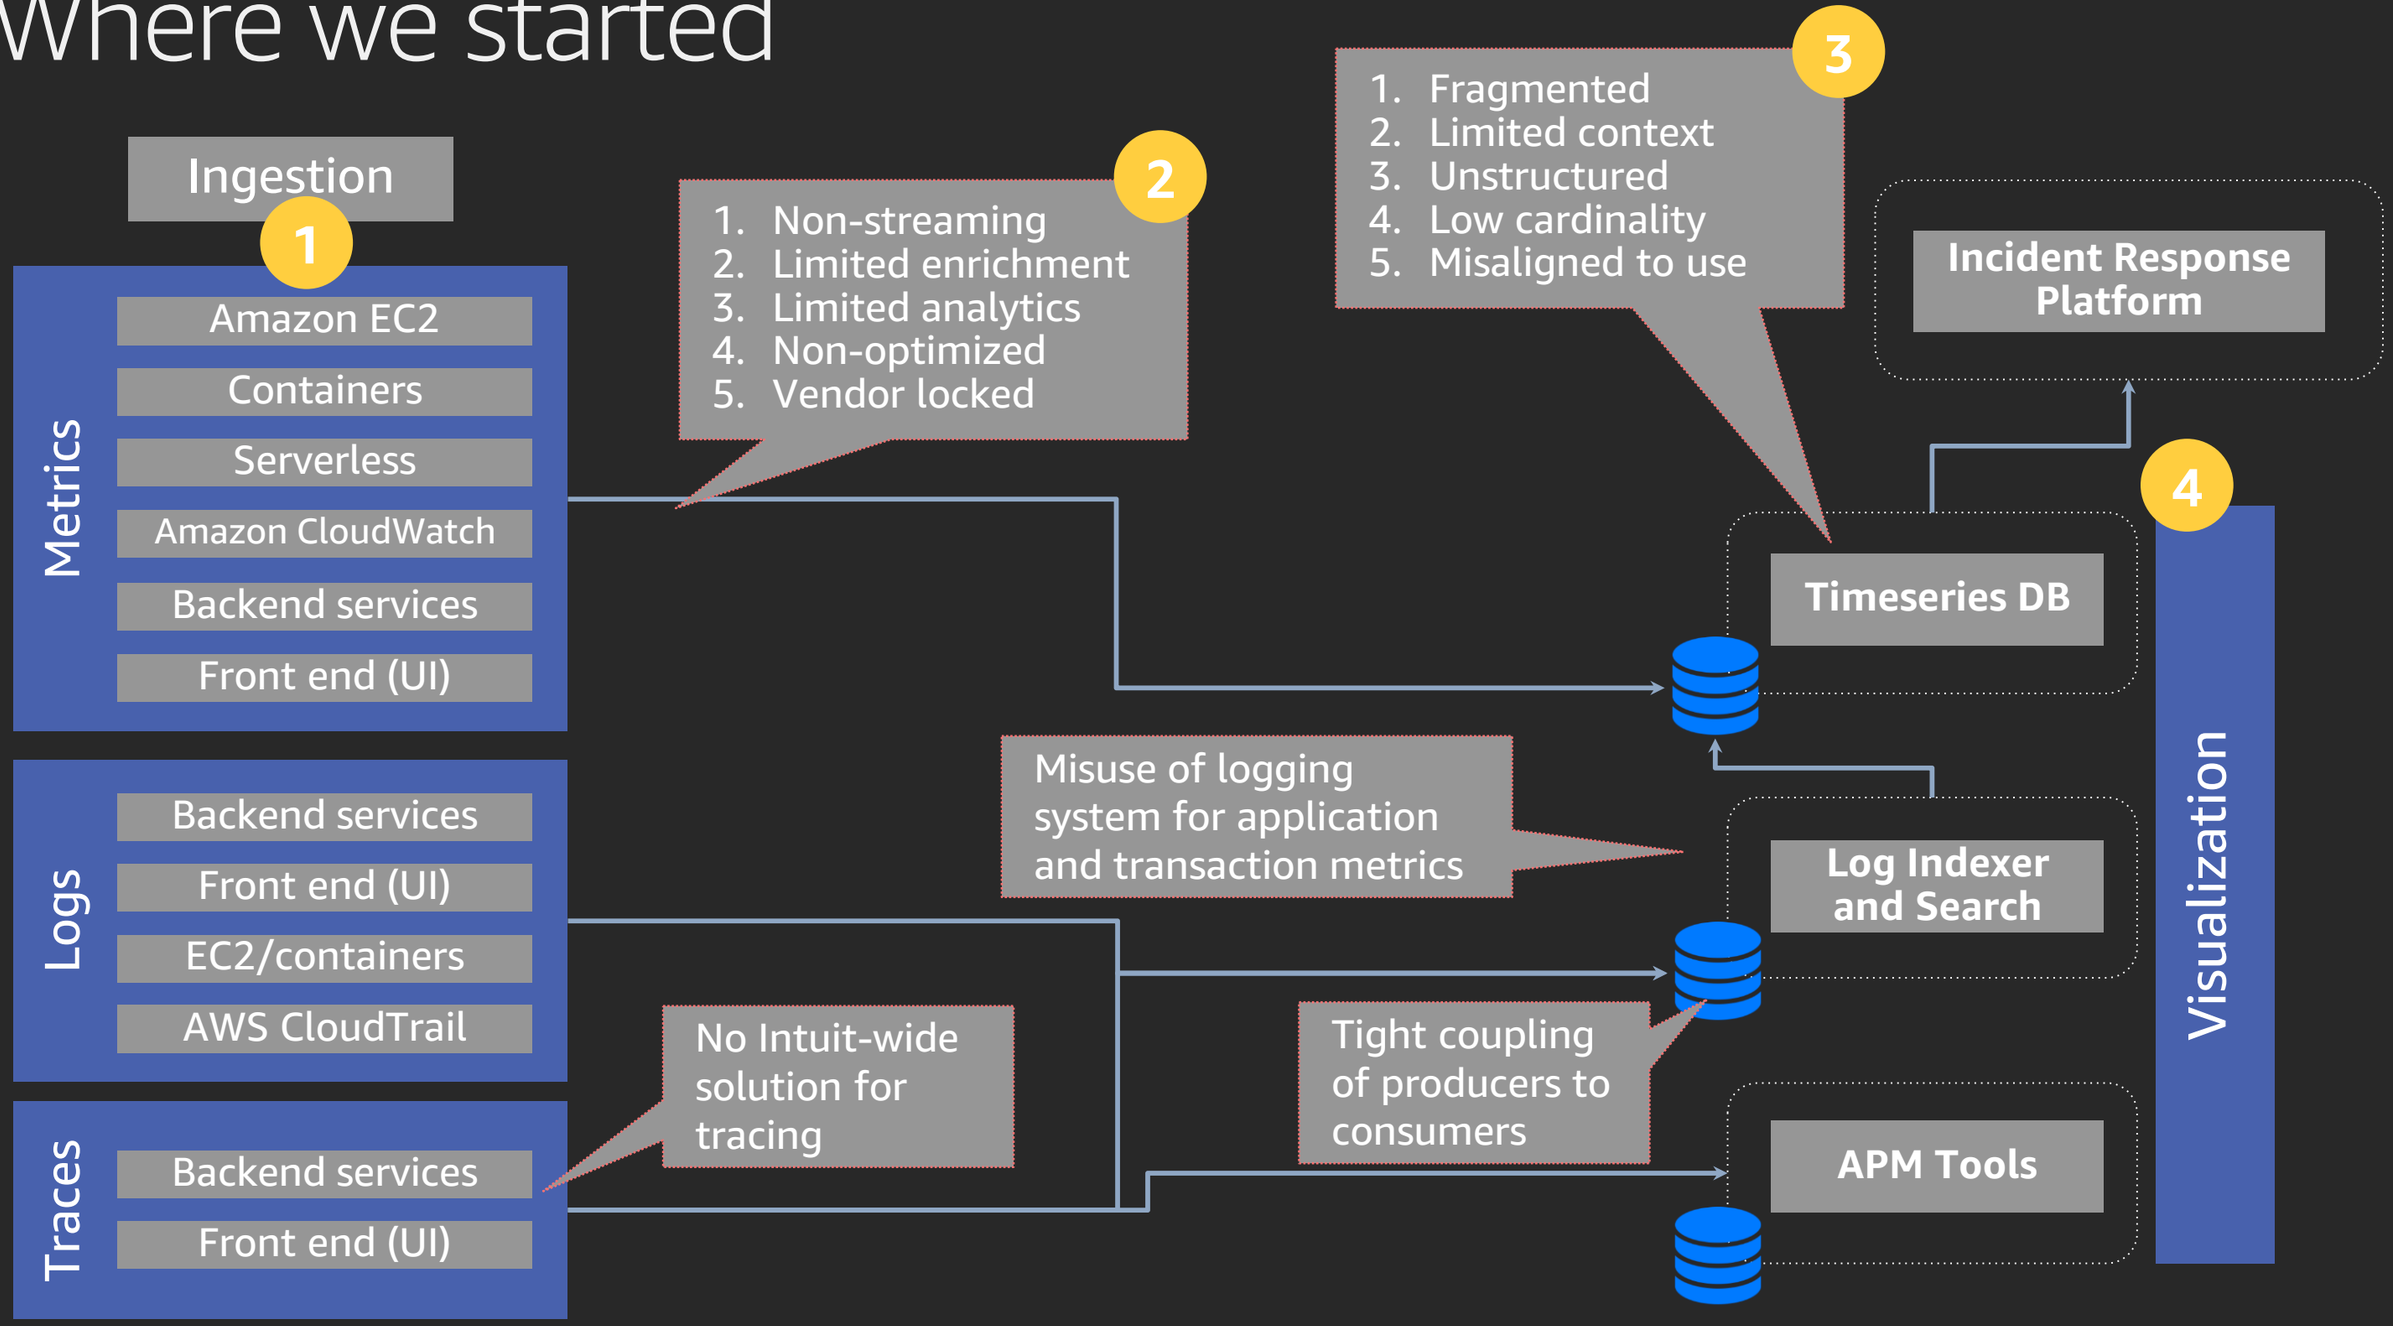
Monitoring took too long to provide data necessary to decide on best path of repair

We rarely identified the right service and type of failure at first shot

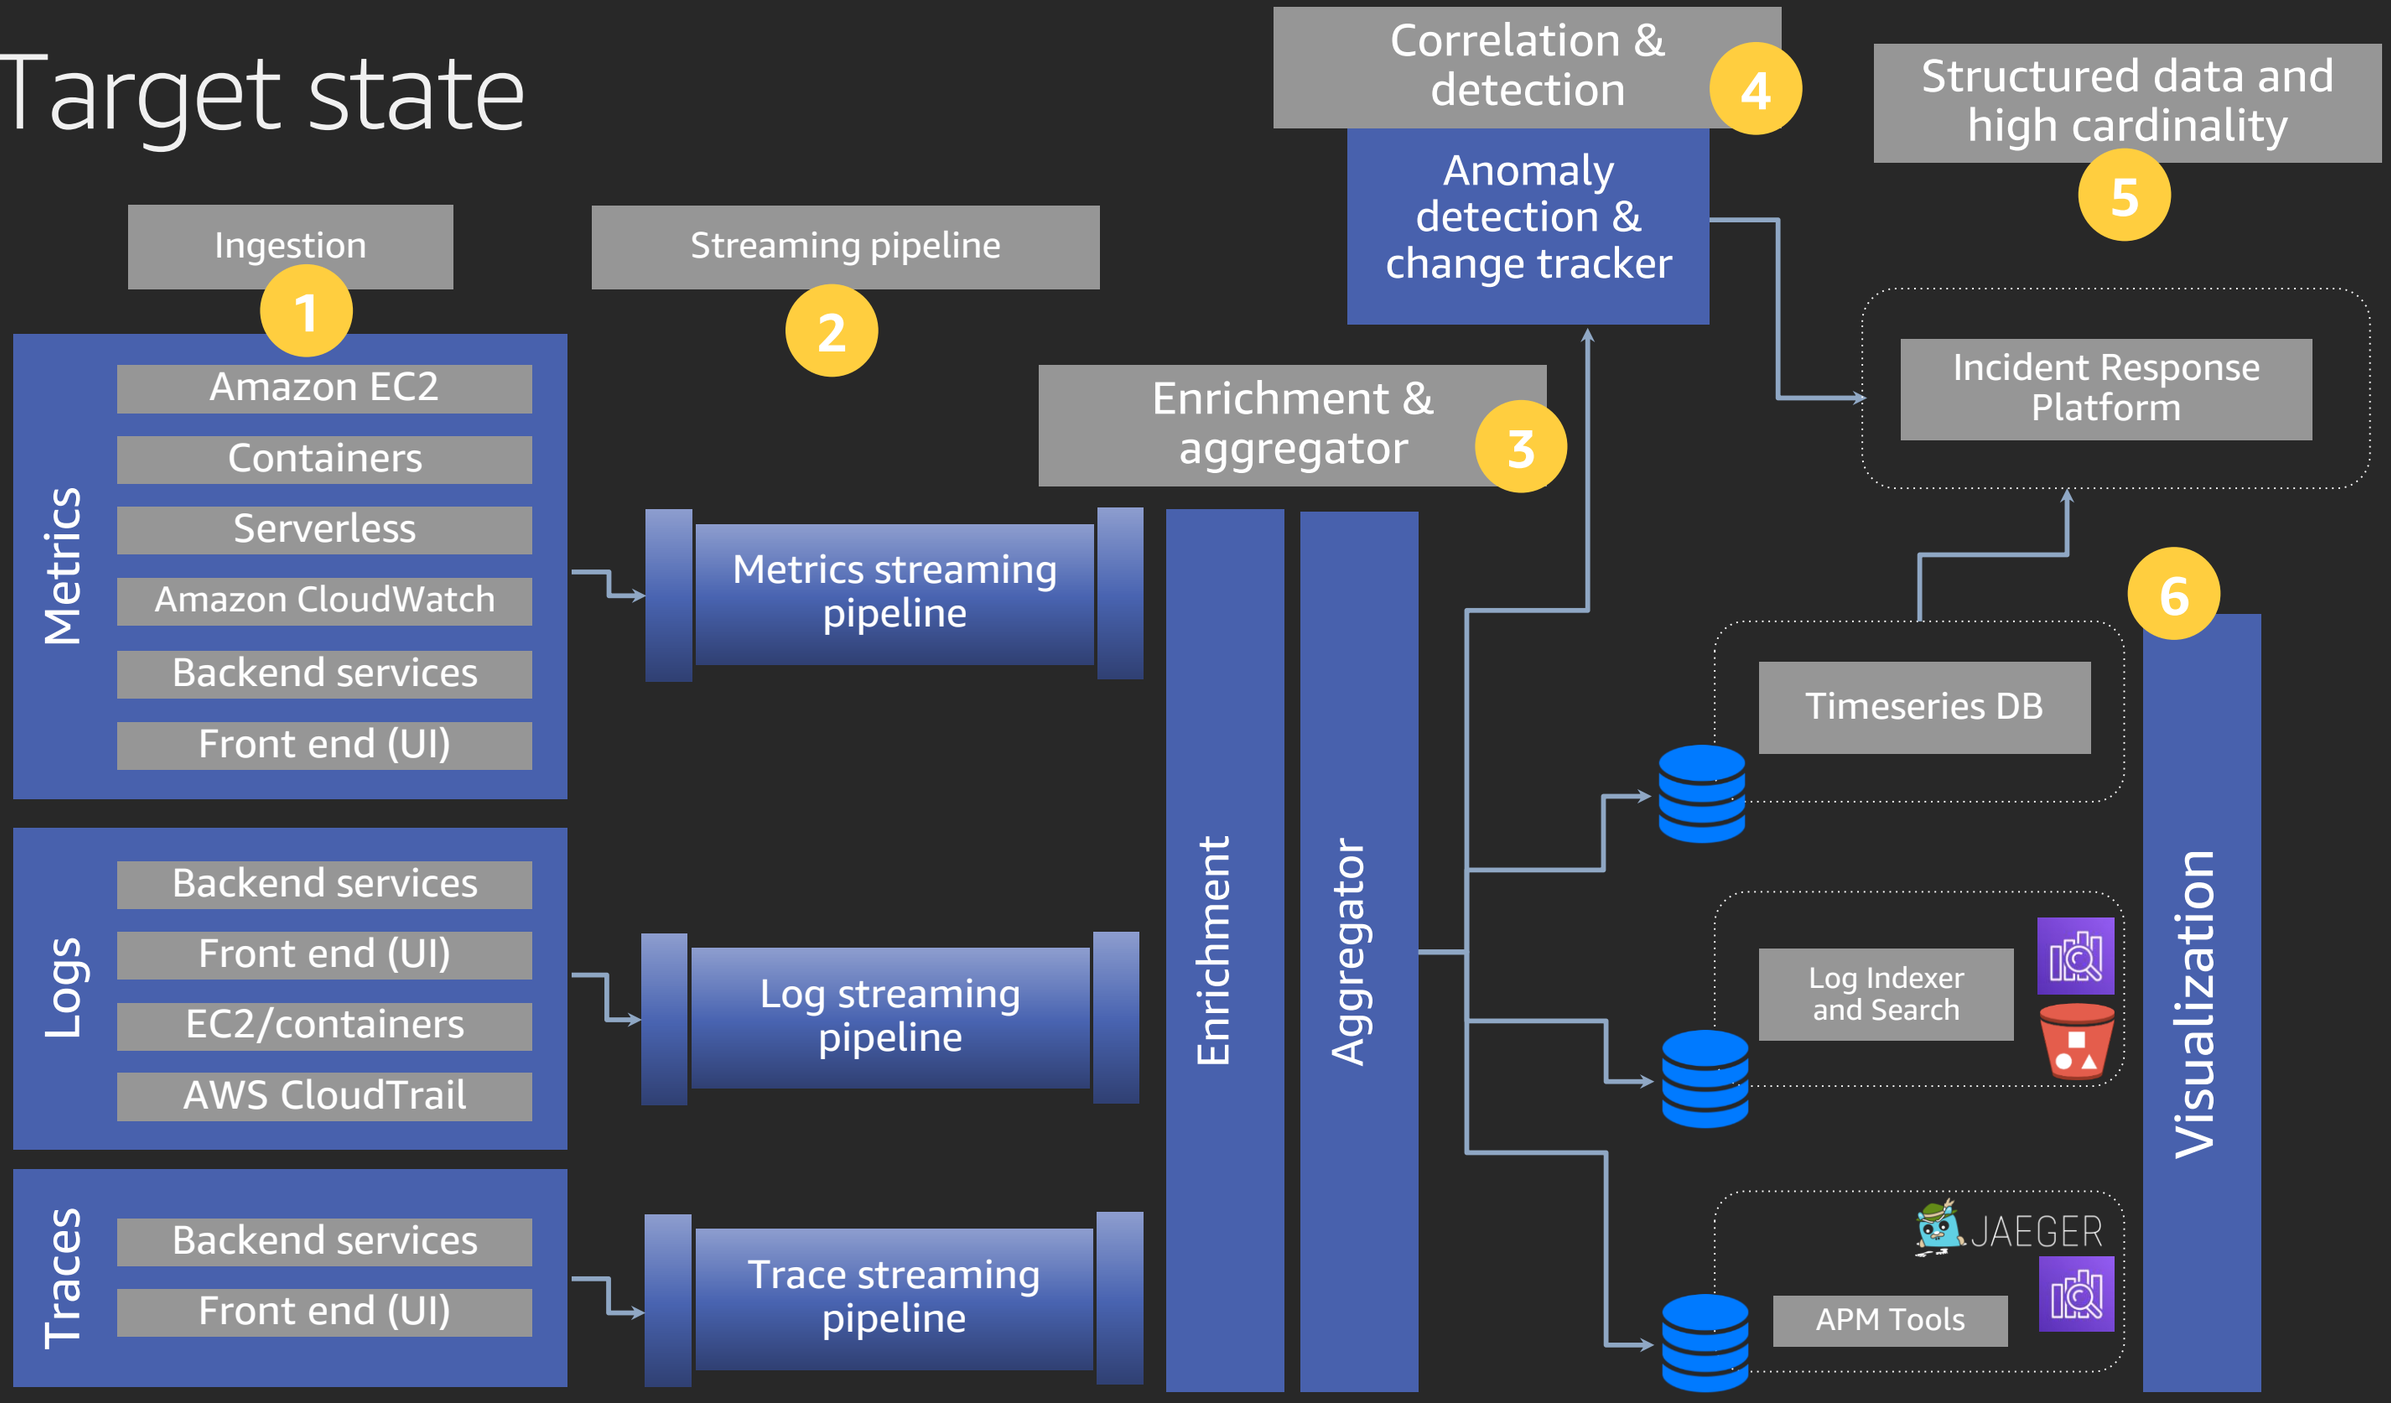
Observability, three pillars

Metrics	Logs	Traces
Numeric data measured at various time intervals (time series data); SLIs (request rate, error rate, duration, CPU%, etc.)	Timestamped records of discrete events that happened within an application or system, such as a failure, an error, or a state transformation	A trace represents a single user's journey across multiple applications and systems (usually microservices)

Where we started



Target state



Tracing

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Distributed Tracing

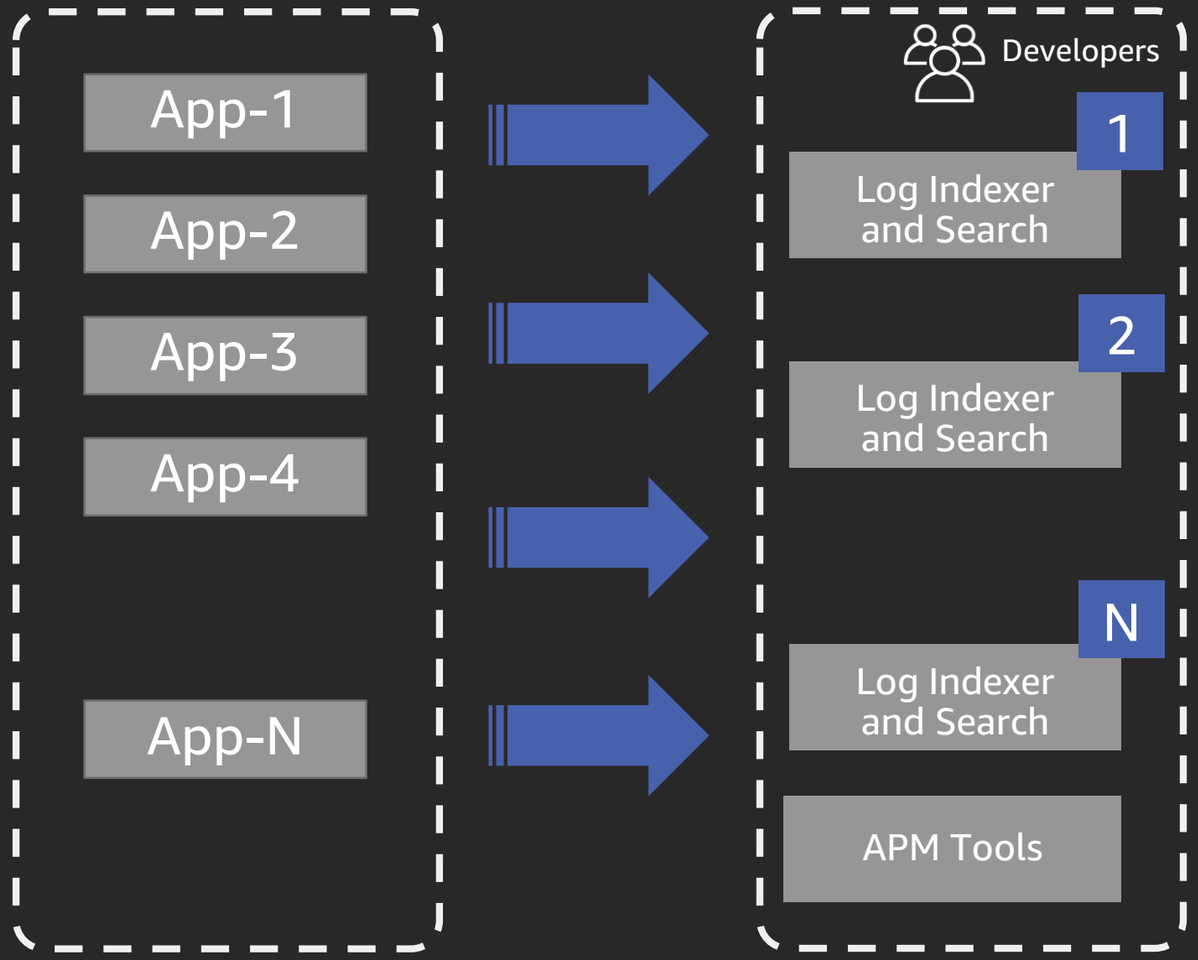
intuit

Process of tracking a request end to end as it traverses through microservices, in near real time

Commonly used terms

- Trace
- Span

Poor man's tracing



Nonstandard trace log format

No single tracing solution across Intuit

APM tools worked great for monolithic apps

Traces (logs) were siloed across many infrastructures

First attempt:

One Intuit Tracing Library

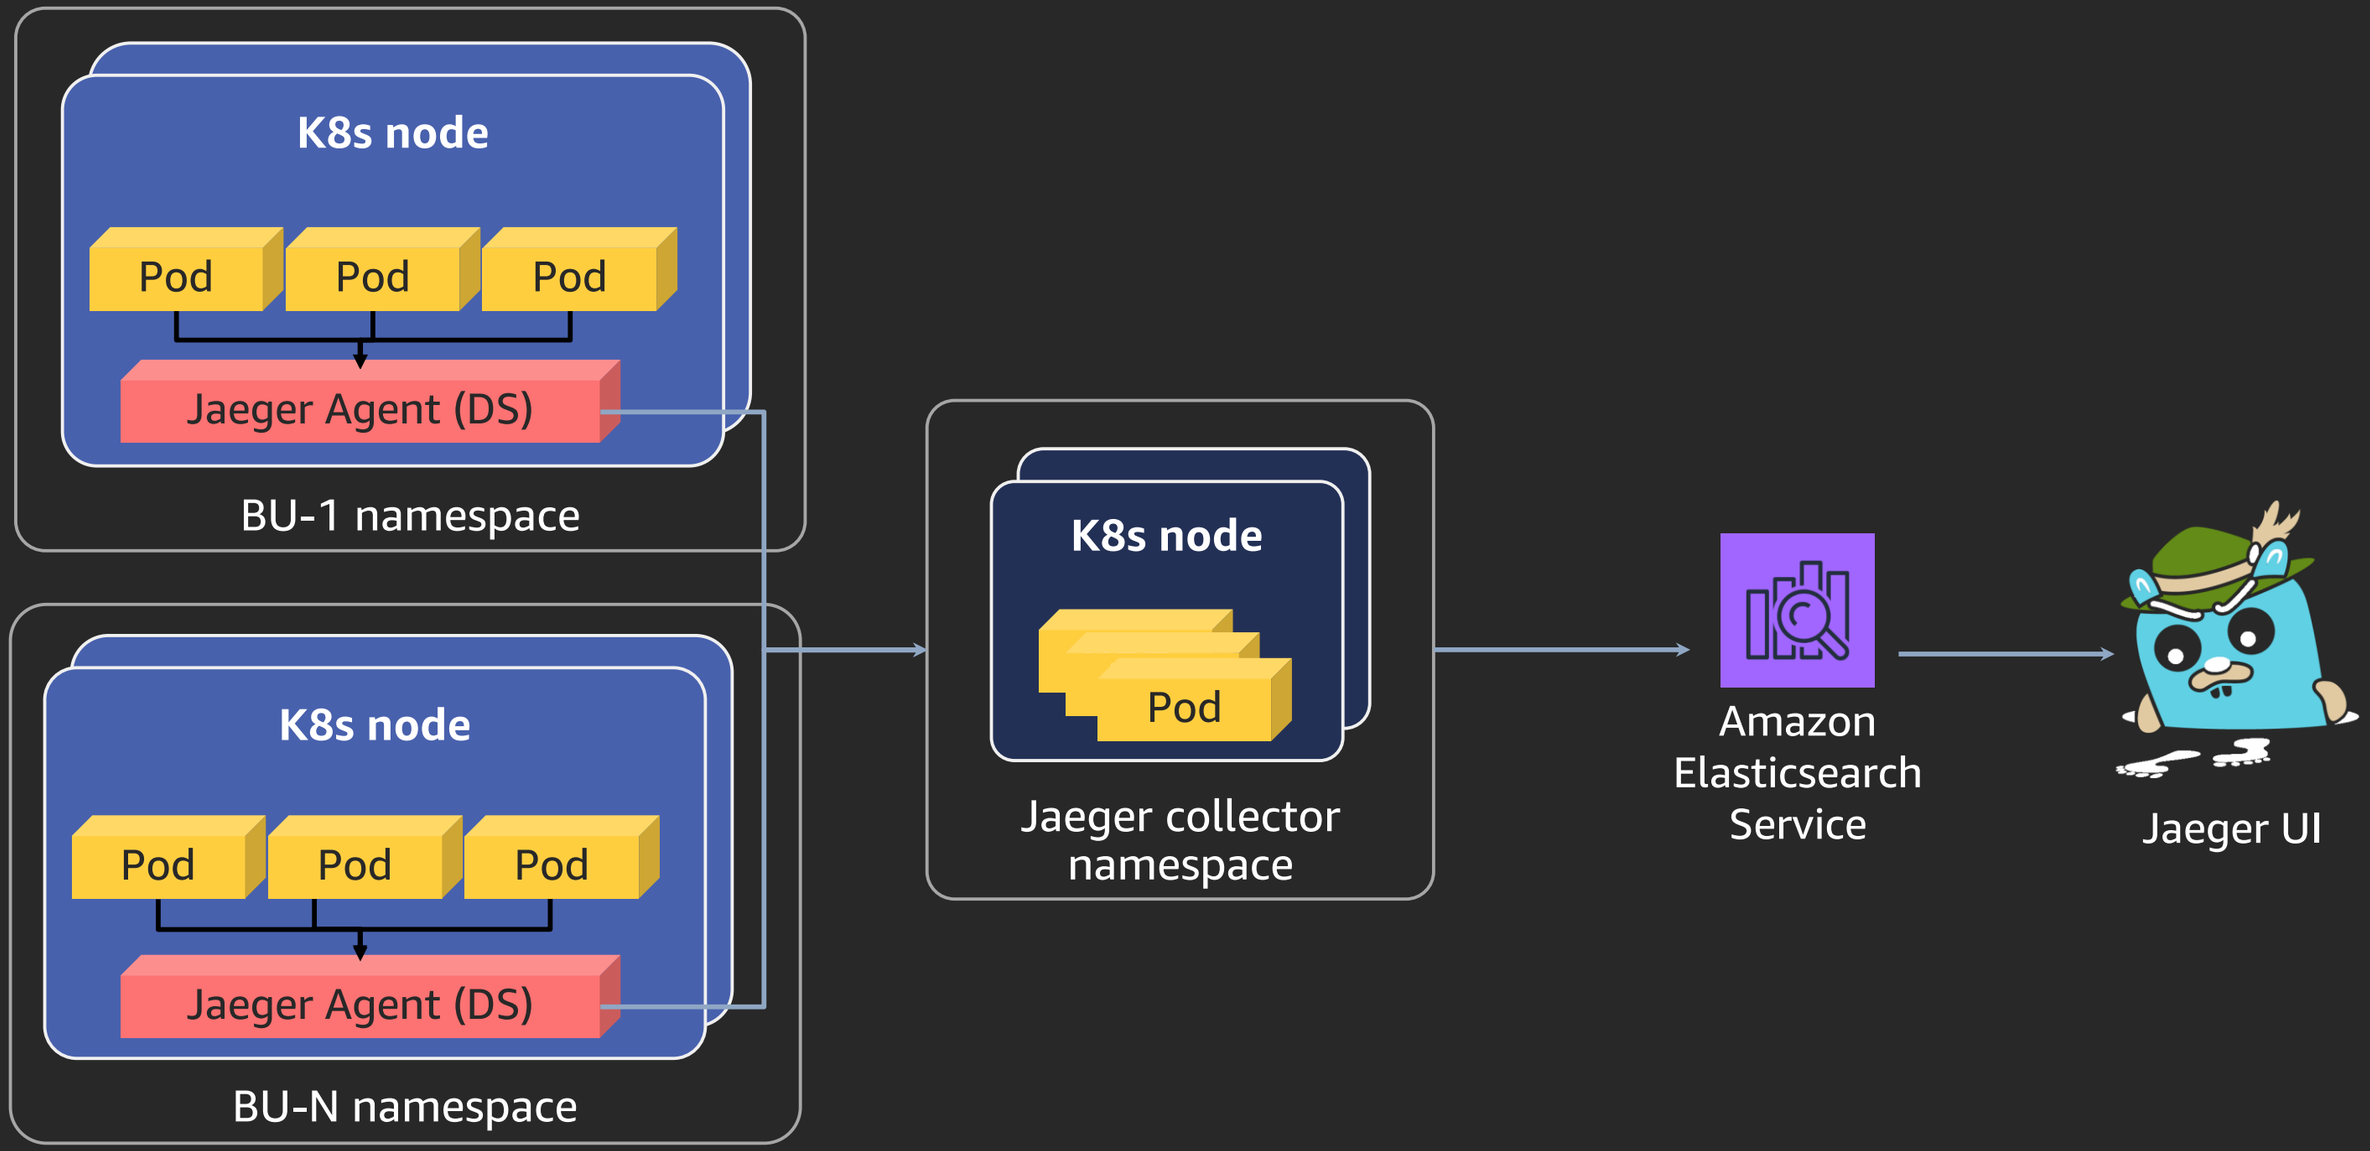
Moving to opentracing standard

Problem – Legacy One Intuit Tracing library didn't follow opentracing standard; impossible to move 100s of services to adopt opentracing library

Solution – Created a new version of library (backward compatible) adhering to opentracing specs

Benefit – Immediate conversion to opentracing for 100s of services by simply flipping the version number

Jaeger with Amazon Elasticsearch Service



Jaegar Collector Configuration

- Flush Interval – 200ms
- Bulk Workers committing to Amazon Elasticsearch Service – 5
- Number of worker pulling from the Jaegar internal queue – 60

Amazon Elasticsearch Service scalability

- Prewarm your cluster
- Refresh interval of 30s (default is 1)
- No more than 1 active shard per node
- Shard Size < 50GB (Recommended)
- Doc Count per shard < 2 billion (Lucene limitation)

Achieved ~60k TPS with 20 nodes on i3.2XL & **~350K TPS** with 40 nodes on i3.16XL instance

Capacity needs without sampling: **2 million** traces per second

Sampling

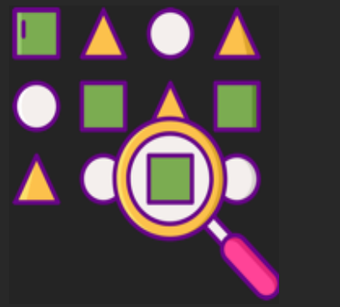
intuit

Head sampling

Root span system makes the decision usually based on probability (e.g., 10%); all other systems in the path of the trace respect the decision

Tail sampling

Sampling decision is made when the full trace information has been gathered



Tracing – Demo

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Metrics

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Types of metrics

intuit

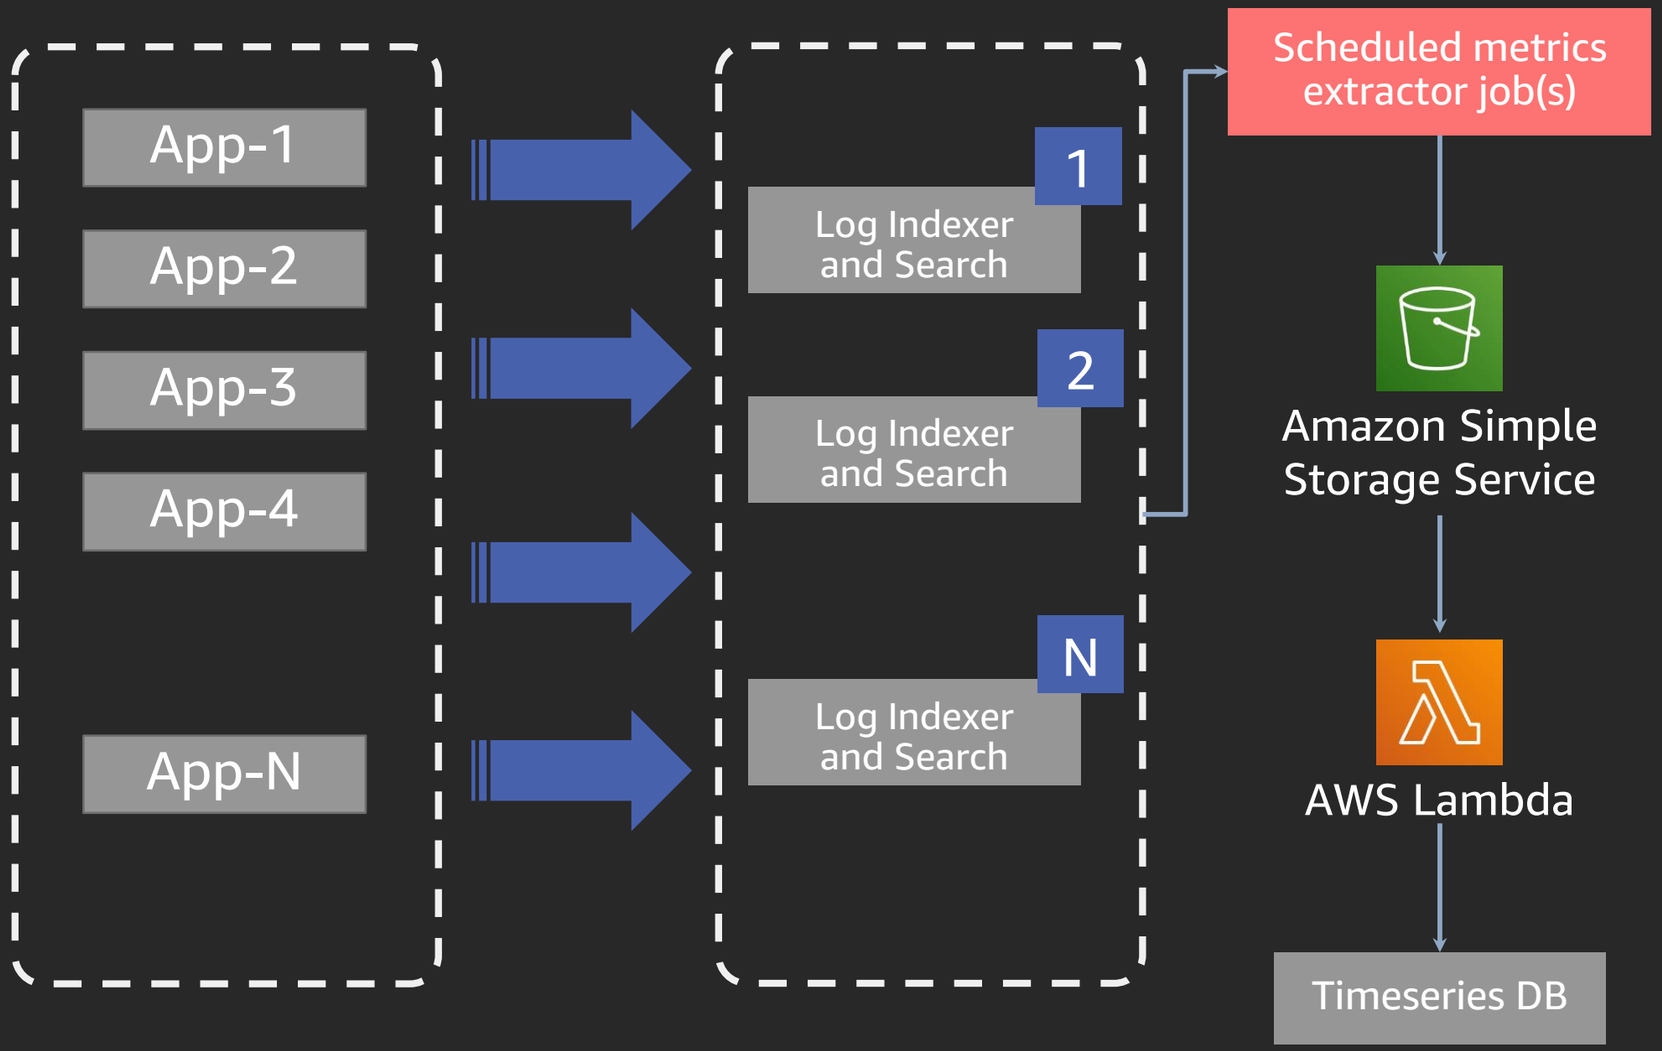
System resource metrics – K8s node/pod CPU, memory metrics

AWS resource metrics (Amazon RDS, Amazon ES, Amazon DynamoDB, etc.) – CloudWatch

Transaction metrics – RED (request rate, error rate, duration)

Real user monitoring (RUM) metrics – failed customer transactions, page load time, etc.

RED metrics – Misusing Logging System



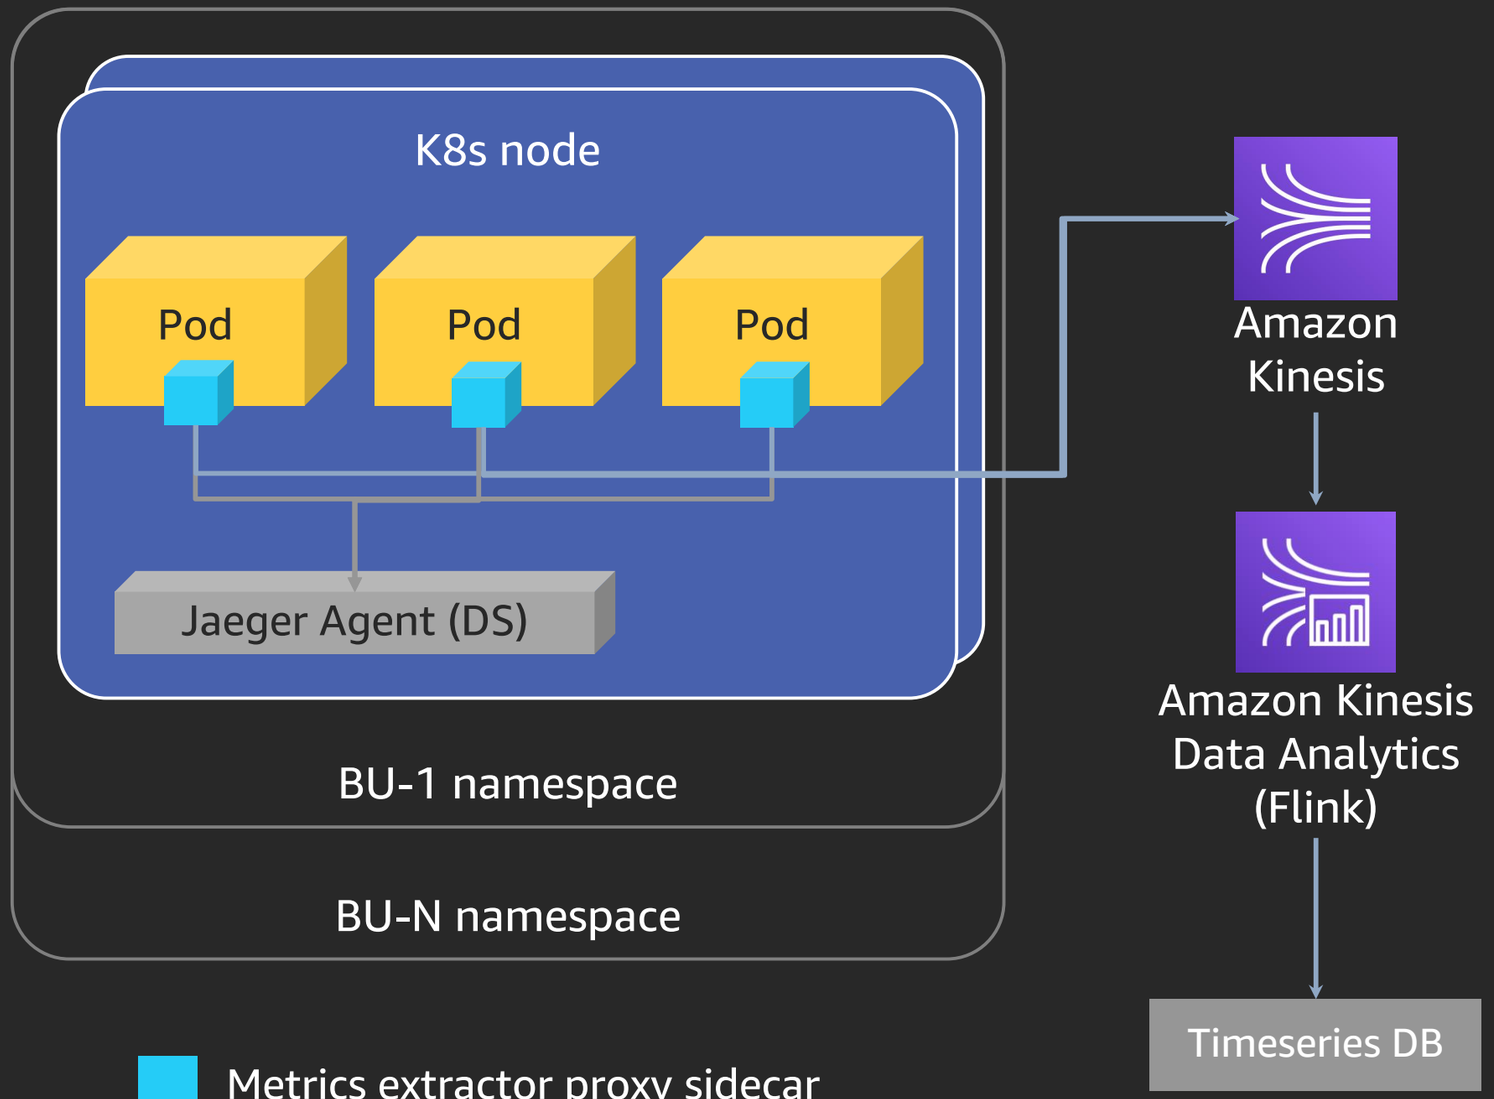
Overloaded Log Indexer and Search

10 minutes of delayed metrics

Skipped jobs

Impacts real-time log searches

RED metrics extraction at source



No additional code instrumentation needed

Batch to real time

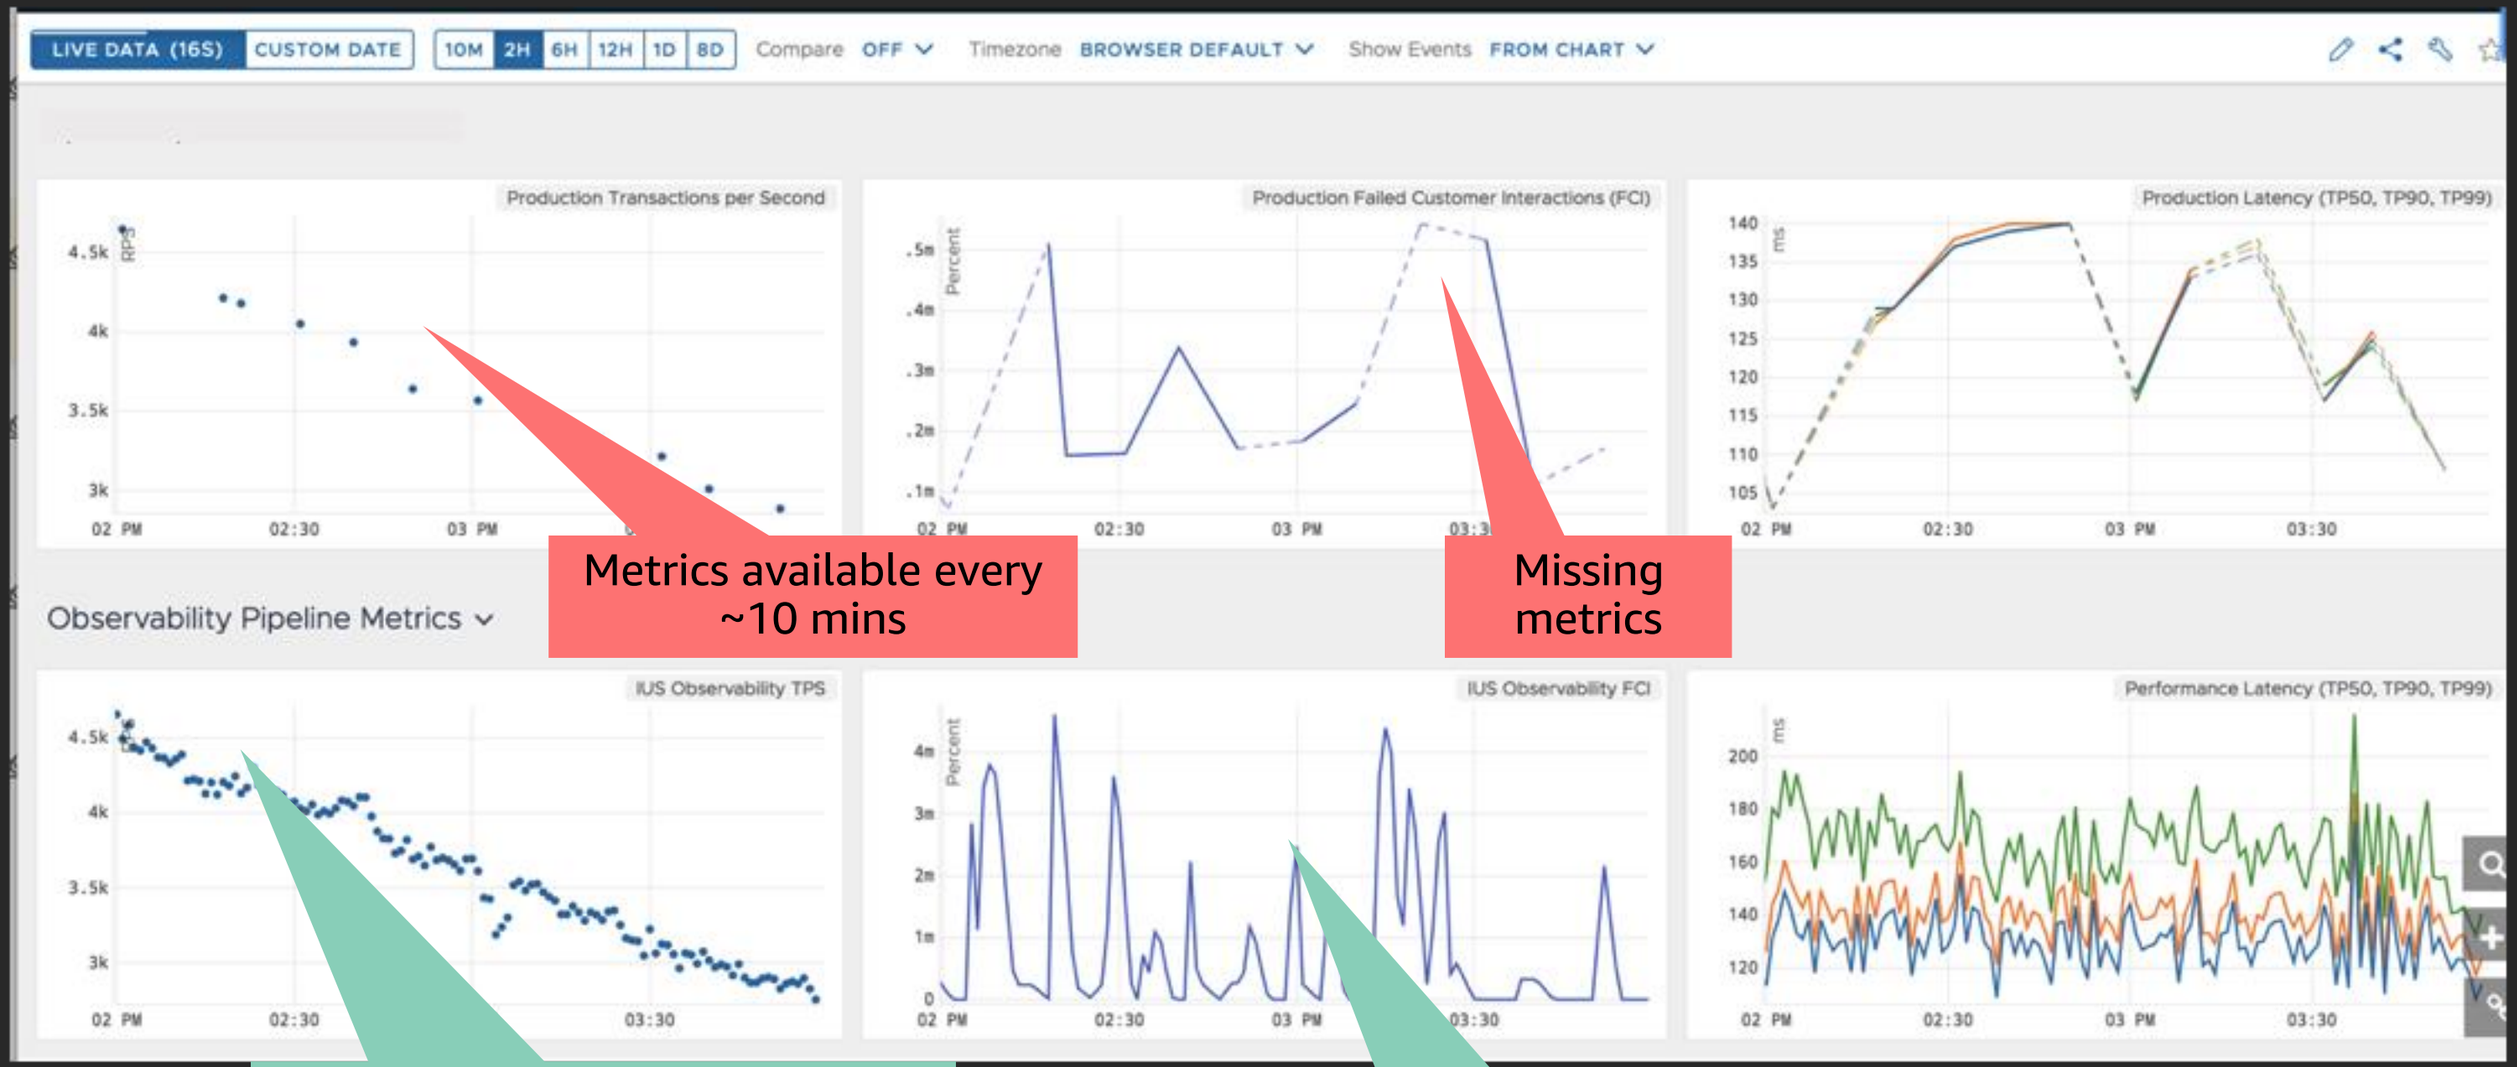
Time to destination for metrics < 30 seconds

AWS Auto Scaling of shard removes capacity issues

Metrics – Demo

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Metrics availability: improved from ~10 mins to <1 min



Log

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Logging in container world

intuit

Problems

Increased log throughput per node in the container world

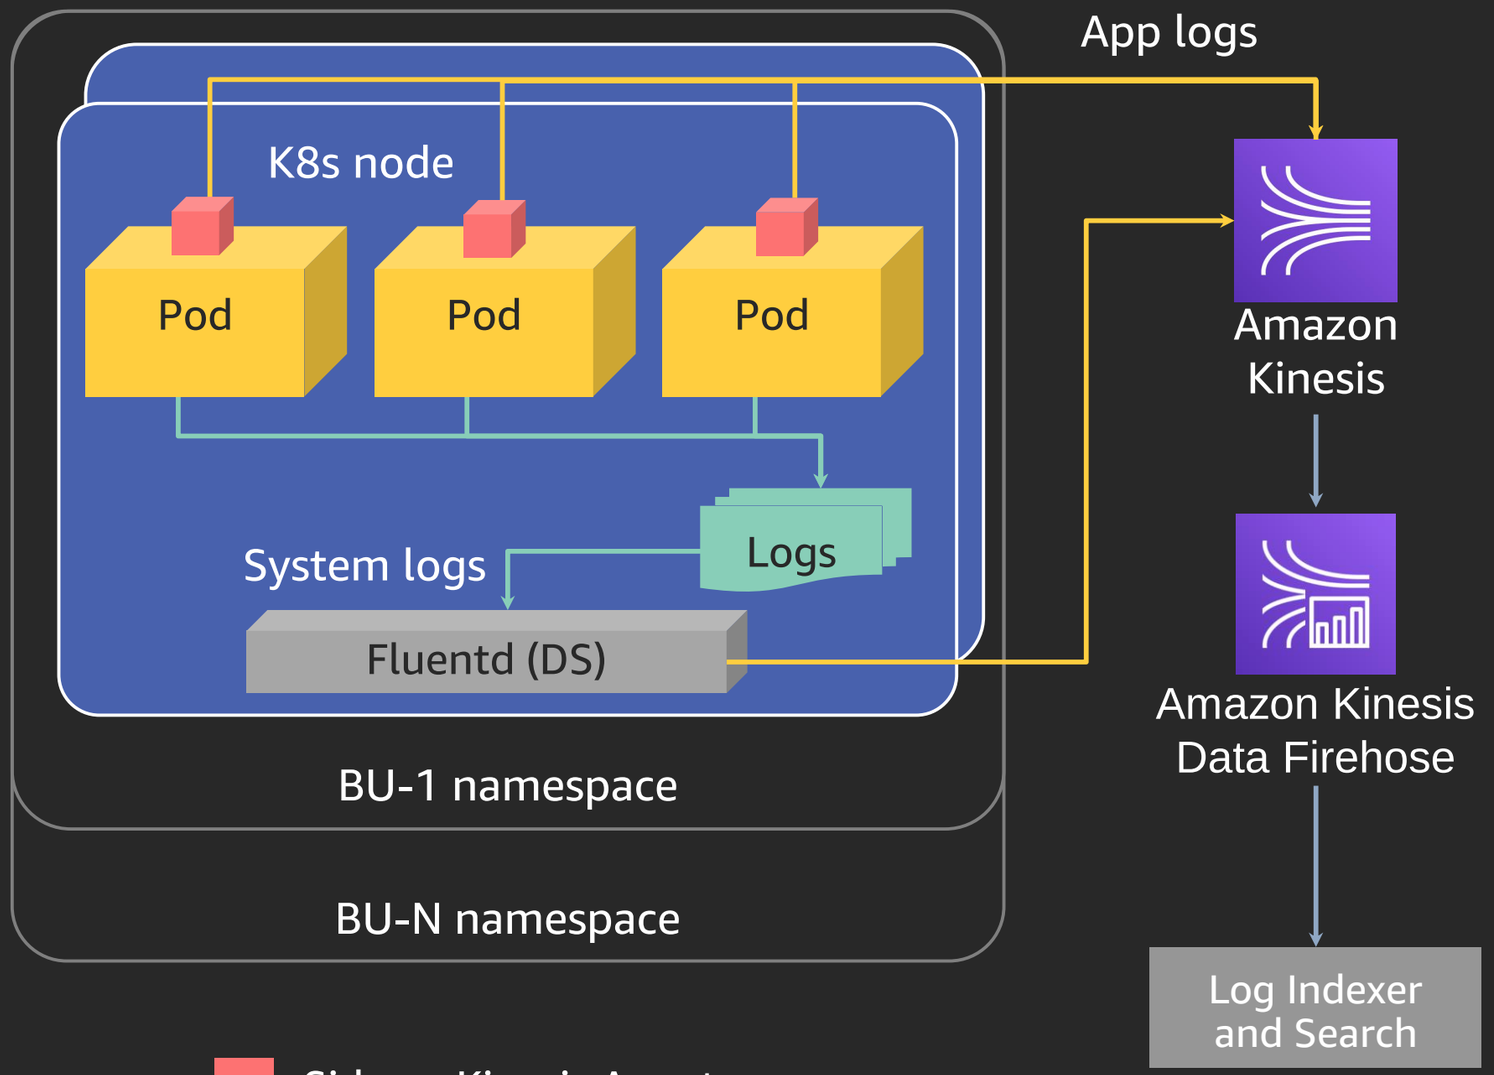
Logs needed additional Kubernetes context

Way too many logs; a major portion of logs never get used

Log losses due to backend capacity constraints

Solution: Highly scalable and highly reliable log streaming pipeline

Log streaming pipeline



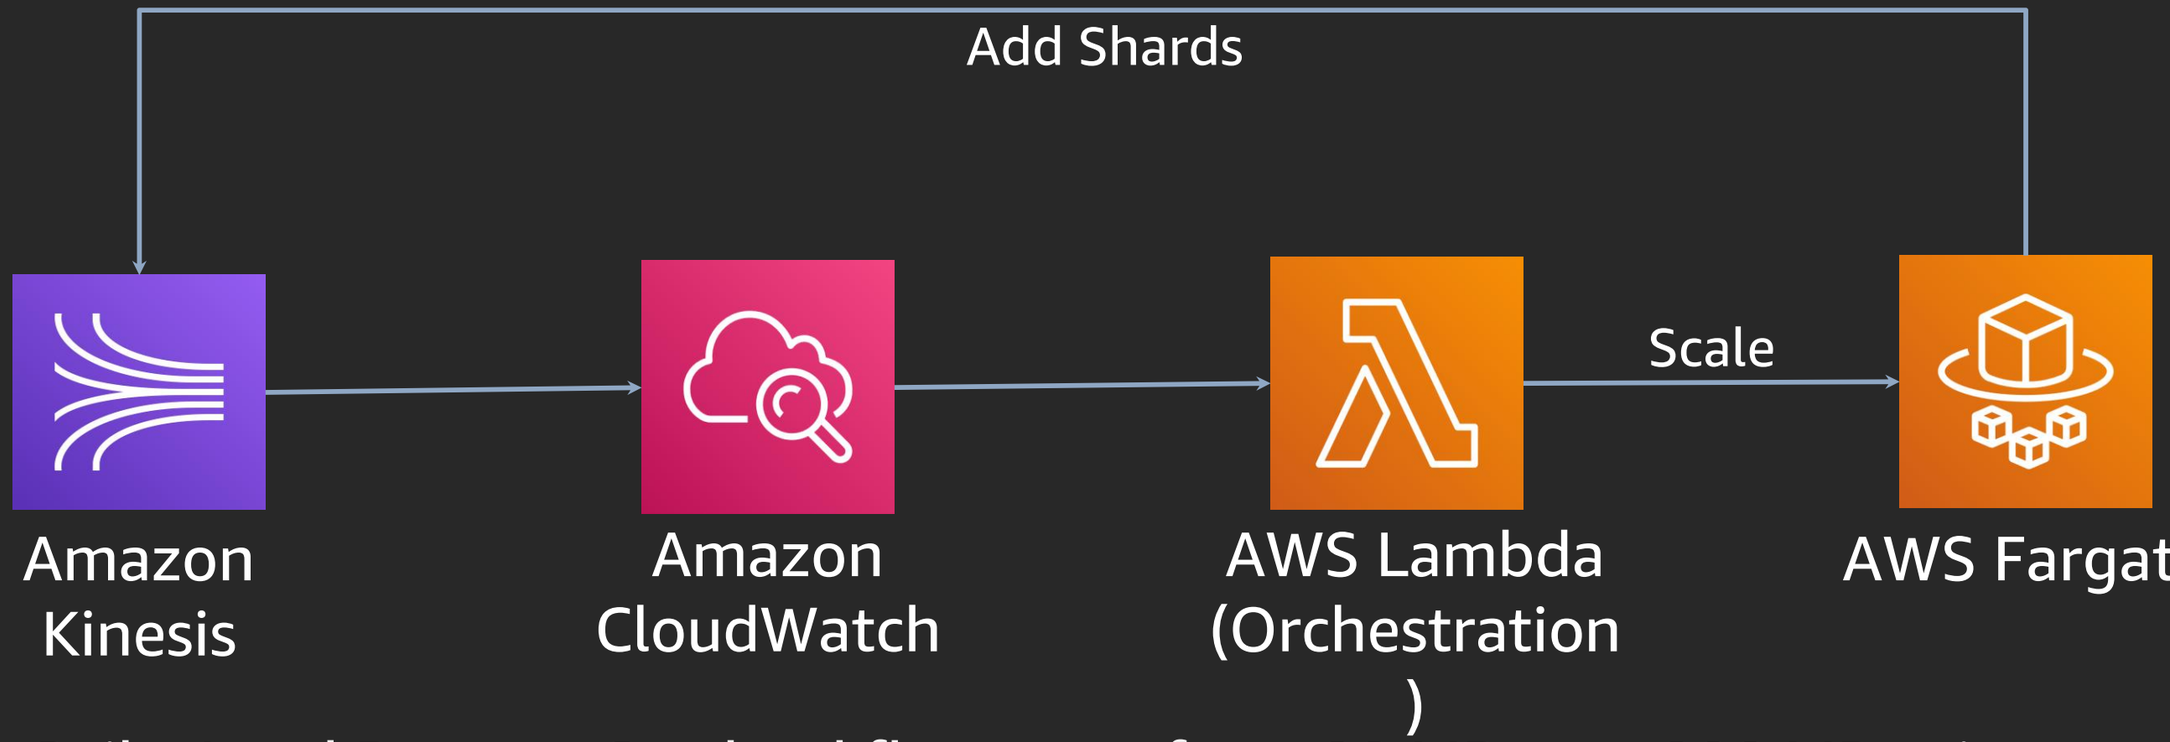
Amazon Kinesis Agent sidecar is able to support high log throughput needs (5000 records per second per pod – 500 milli core)

FluentD Daemonset is able to support 2500 records per second per node – 700 milli core.

Streaming pipeline opens up opportunities for log routing, filtering, and transformations

Possibility for various log backend (depending on use cases)

Amazon Kinesis auto scaling



- Spiky Load Pattern: Our load fluctuates from 80 MBps to 800 MBps in a week
- Triggers scaling when load breaches 70% of ingestion rate or in case of a Hot Shard
- Faster to scale in whole number multiples compared to fraction (10 -> 20 takes 3.5m but 10→ 15 takes 4.5m)

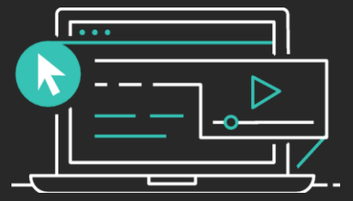
Reference

intuit

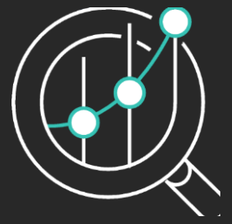
https://tinyurl.com/2019ANT330

Learn big data with AWS Training and Certification

Resources created by the experts at AWS to help you build and validate data analytics skills



New free digital course, Data Analytics Fundamentals, introduces Amazon S3, Amazon Kinesis, Amazon EMR, AWS Glue, and Amazon Redshift



Classroom offerings, including Big Data on AWS, feature AWS expert instructors and hands-on labs



Validate expertise with the **AWS Certified Big Data** - **Specialty** exam or the new **AWS Certified Data Analytics** - **Specialty** beta exam

Visit aws.amazon.com/training/paths-specialty/

Thank you!

Balaji Arunachalam & **Nagaraja Tantry**

balaji_arunachalam@intuit.com
nagaraja_tantry@intuit.com

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws



Please complete the session survey in the mobile app.

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.