Distributed Tracing on ClickHouse

eBay Inc.

OSACON-2021

November 2nd, 2021

Sudeep Kumar, Amber Vaidya

OSACON 2021

eBay Scale

Source: Euromonitor International, 2017; Note: Regions are North America, Latin America, Europe, Asia-Pacific and Australalia/New Zealand. Latin America includes Brazil and Mexico.

- Global Commerce Leader with presence in more than 190 markets
- 154 Million Active Buyers
- $19.5 Billion GMV in Q3 2021
- $2.5 Billion in Revenue in Q3 2021

OSA CON 2021

Sherlock.io

- Observability platform for EBAY's monitoring needs
- Supports different telemetry signals
  - *Metrics*
  - *Logs*
  - *Events*
  - *Traces*
- Scale of Operation (millions per second)

Platform Overview

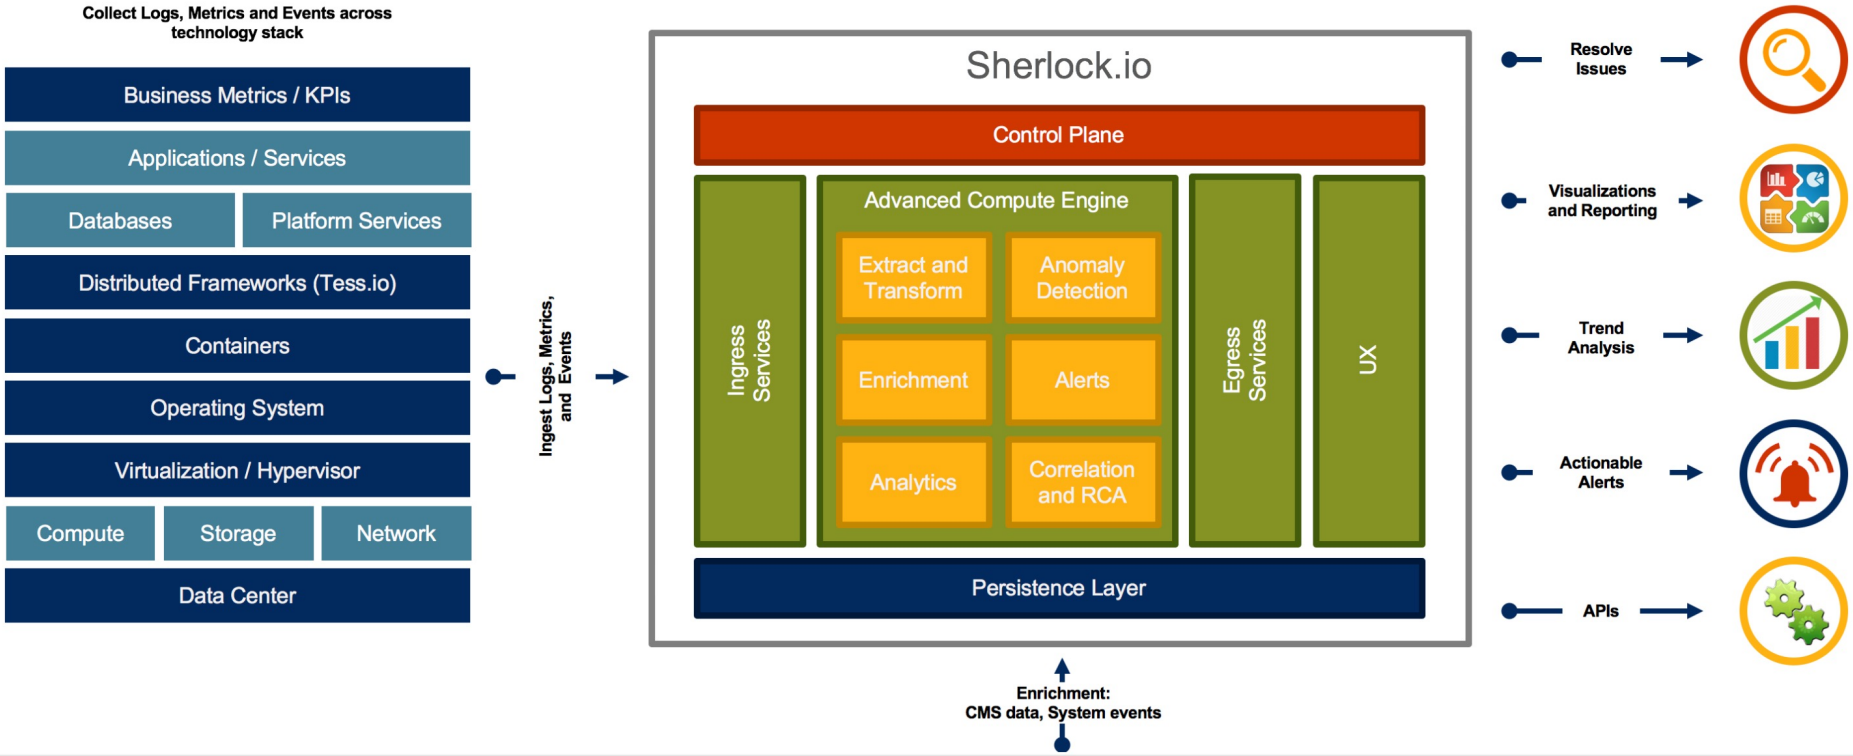

Tracing

- Monitoring signals
  - Metrics
  - Logs/Events
  - Traces
- Distributed Tracing
  - Tracks request flows through multiple applications
  - Helps identify errors/bugs, performance issues
  - Service dependencies

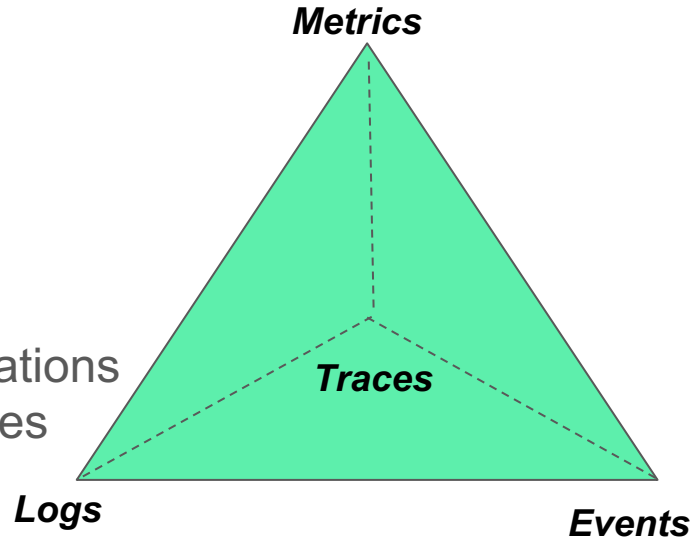

Use-case: I know my 99% percentile latency of my service is 10s, but in what context was high latency observed. Also, was there an event or log associated with that context ?

Open Telemetry (OTEL)

- Observability framework
  - CNCF project
  - Provides agents/libraries
- OTEL Collector
  - vendor-agnostic implementation
  - receive, process & export telemetry data

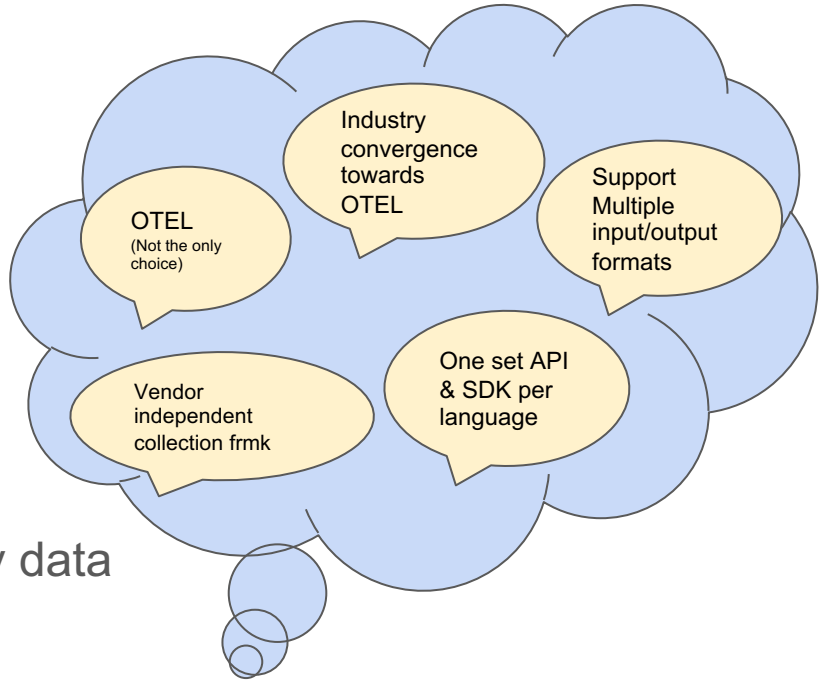

OTEL Collector

- Consists of receivers/exporters
  - Supports various i/o formats
- Intermediate Processors
  - Scrubbing
  - Metadata decorators (k8s processors)
  - RL/Sampling
- Exports to a custom ingestion endpoint
  - Otlp/Otlphttp format

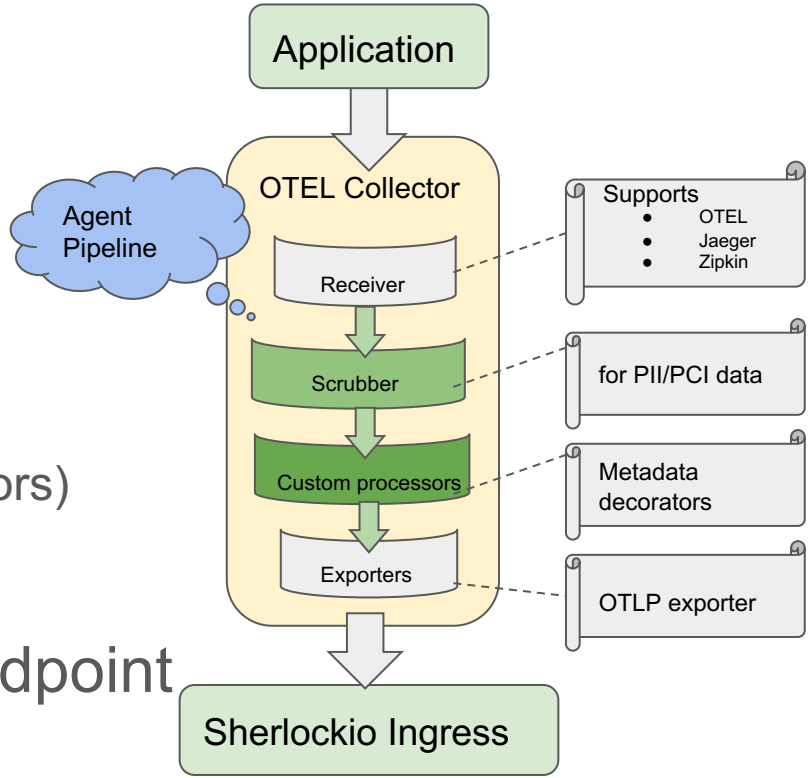

Visualization

- Egress compliance with Jaeger Query APIs
- UX experience
  - Grafana jaeger plugin
  - Optional Jaeger UI

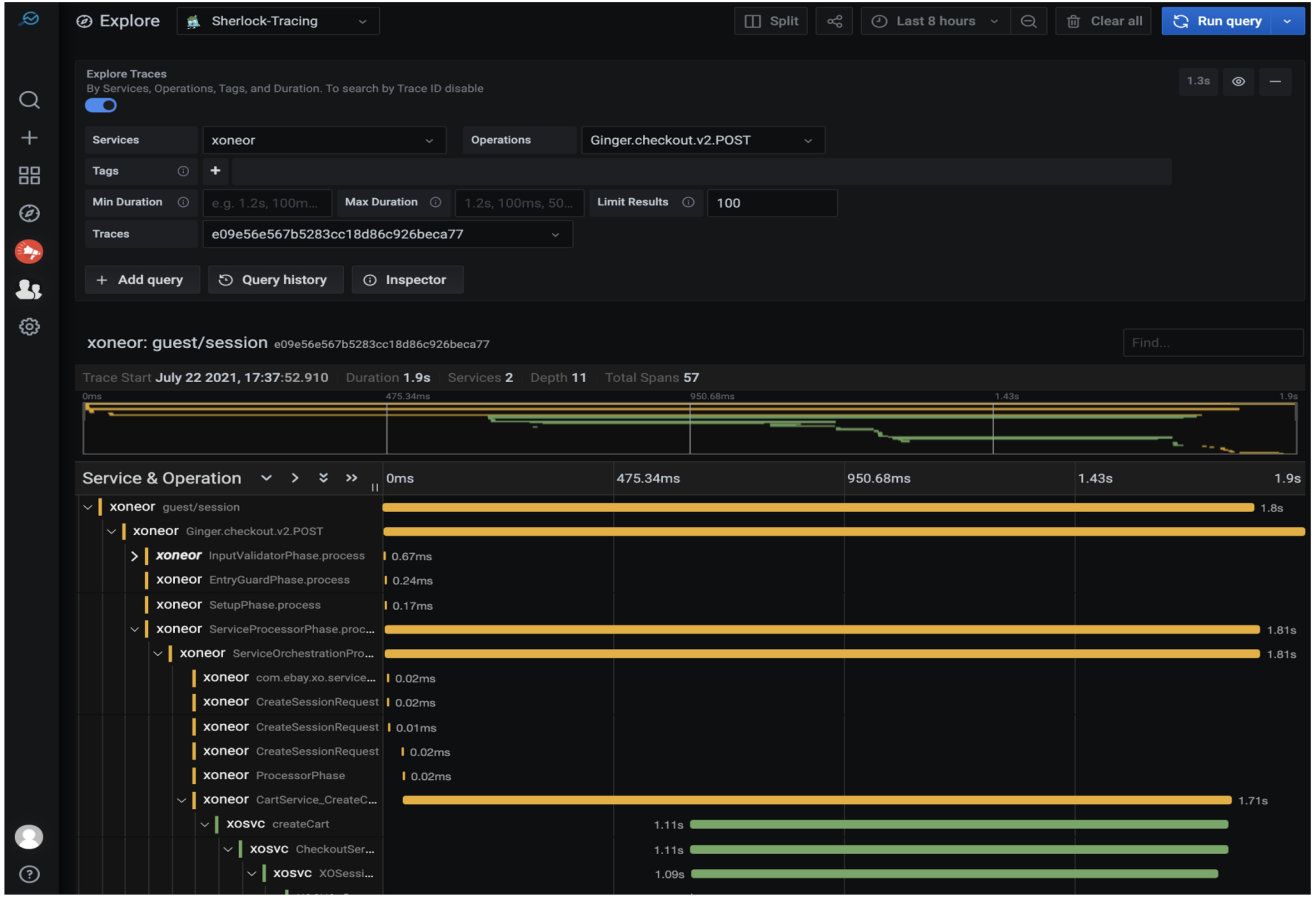

OSA CON 2021

Storage evaluations

- Storage Options

- Criteria Used
  - High ingestion volumes
  - Can support Jaeger Query API Patterns
  - Point Trace-Id lookups for a time interval

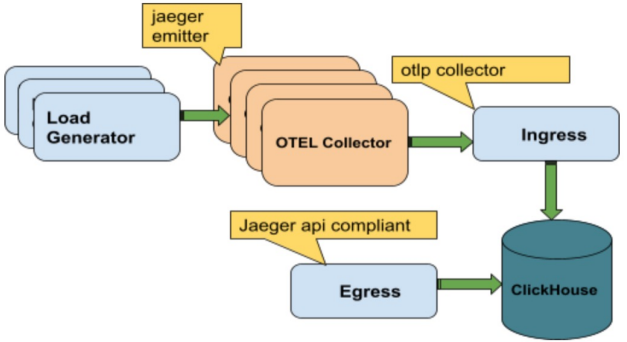

Some Observations

- Ingestion
  - 350K/sec - 4 VPU, 16GB
  - 3KB Span size – 20 Billion rows – 300GB
  - Avg Tag attribute length - 20
- Query patterns evaluated
  - Get Services/Operations
  - Get spans for given a service & TraceId
  - Get by TraceId
  - Search by Attributes (slowest) < 10 QPS

Why ClickHouse ?

- Open source & works well with structured data
- Works well for high ingestion use-cases
- Storage footprint is significantly lower
- Kubernetes operator support
- Extensive infrastructure monitoring in place
- Good inhouse ClickHouse expertise

Deployment

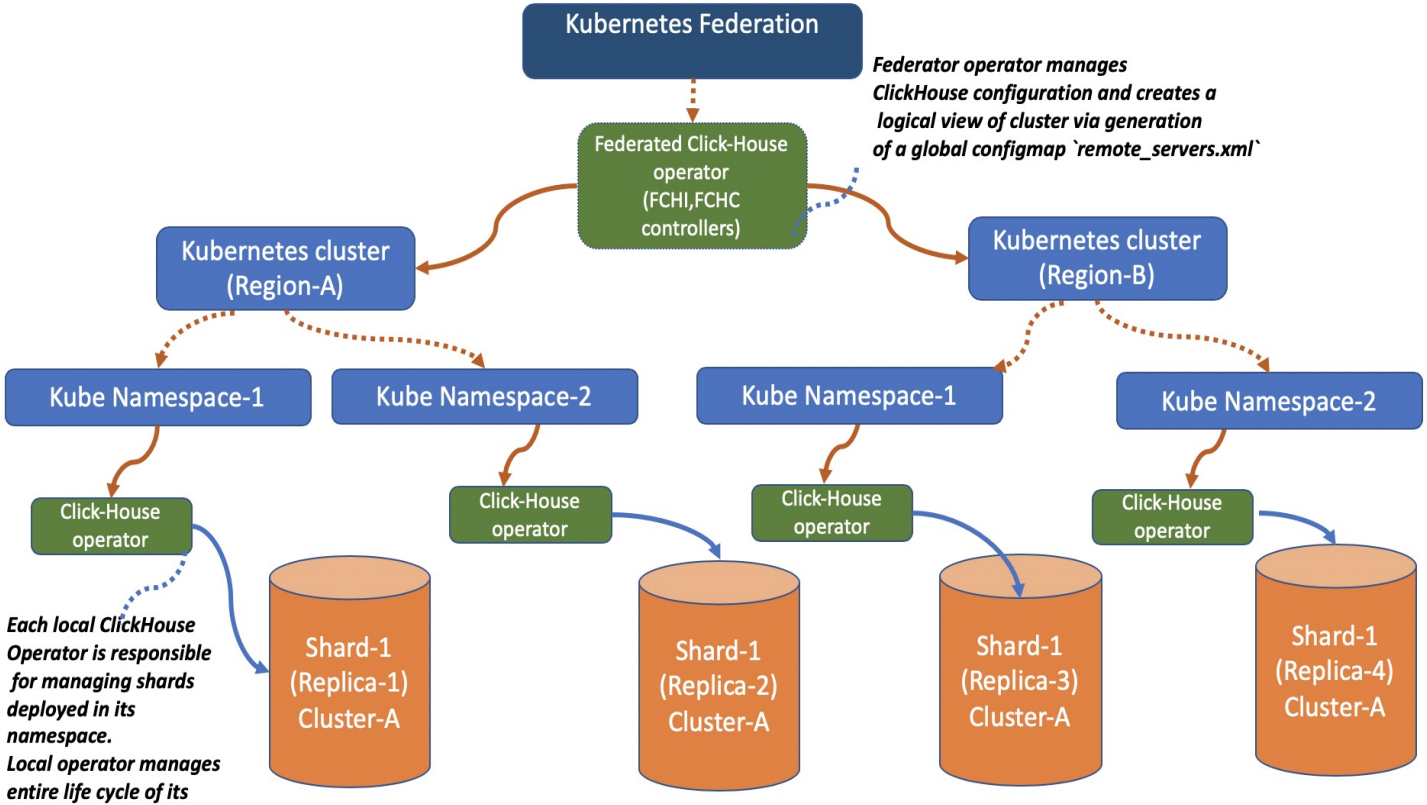
Kubernetes Federation
Federator operator manages ClickHouse configuration and creates a logical view of cluster via generation of a global configmap `remote_servers.xml`
Federated Click-House operator (FCHI,FCHC controllers)
Kubernetes cluster (Region-A)
Kubernetes cluster (Region-B)
Kube Namespace-1
Kube Namespace-2
Kube Namespace-1
Kube Namespace-2
Click-House operator
Click-House operator
Click-House operator
Click-House operator
Each local ClickHouse Operator is responsible for managing shards deployed in its namespace.
Local operator manages entire life cycle of its responsible shards
Shard-1 (Replica-1) Cluster-A
Shard-1 (Replica-2) Cluster-A
Shard-1 (Replica-3) Cluster-A
Shard-1 (Replica-4) Cluster-A

Deployment Contd..

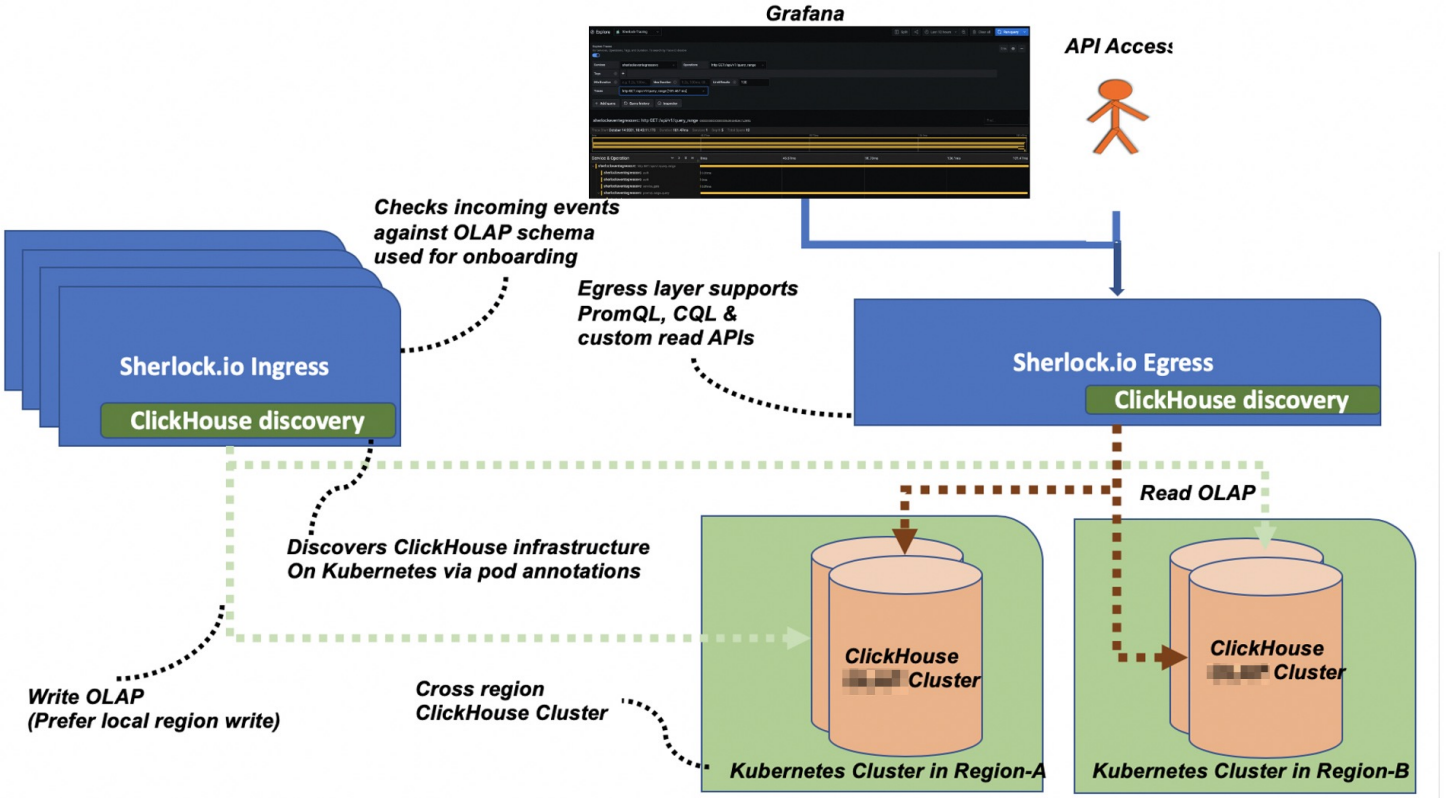

Tracing pipeline

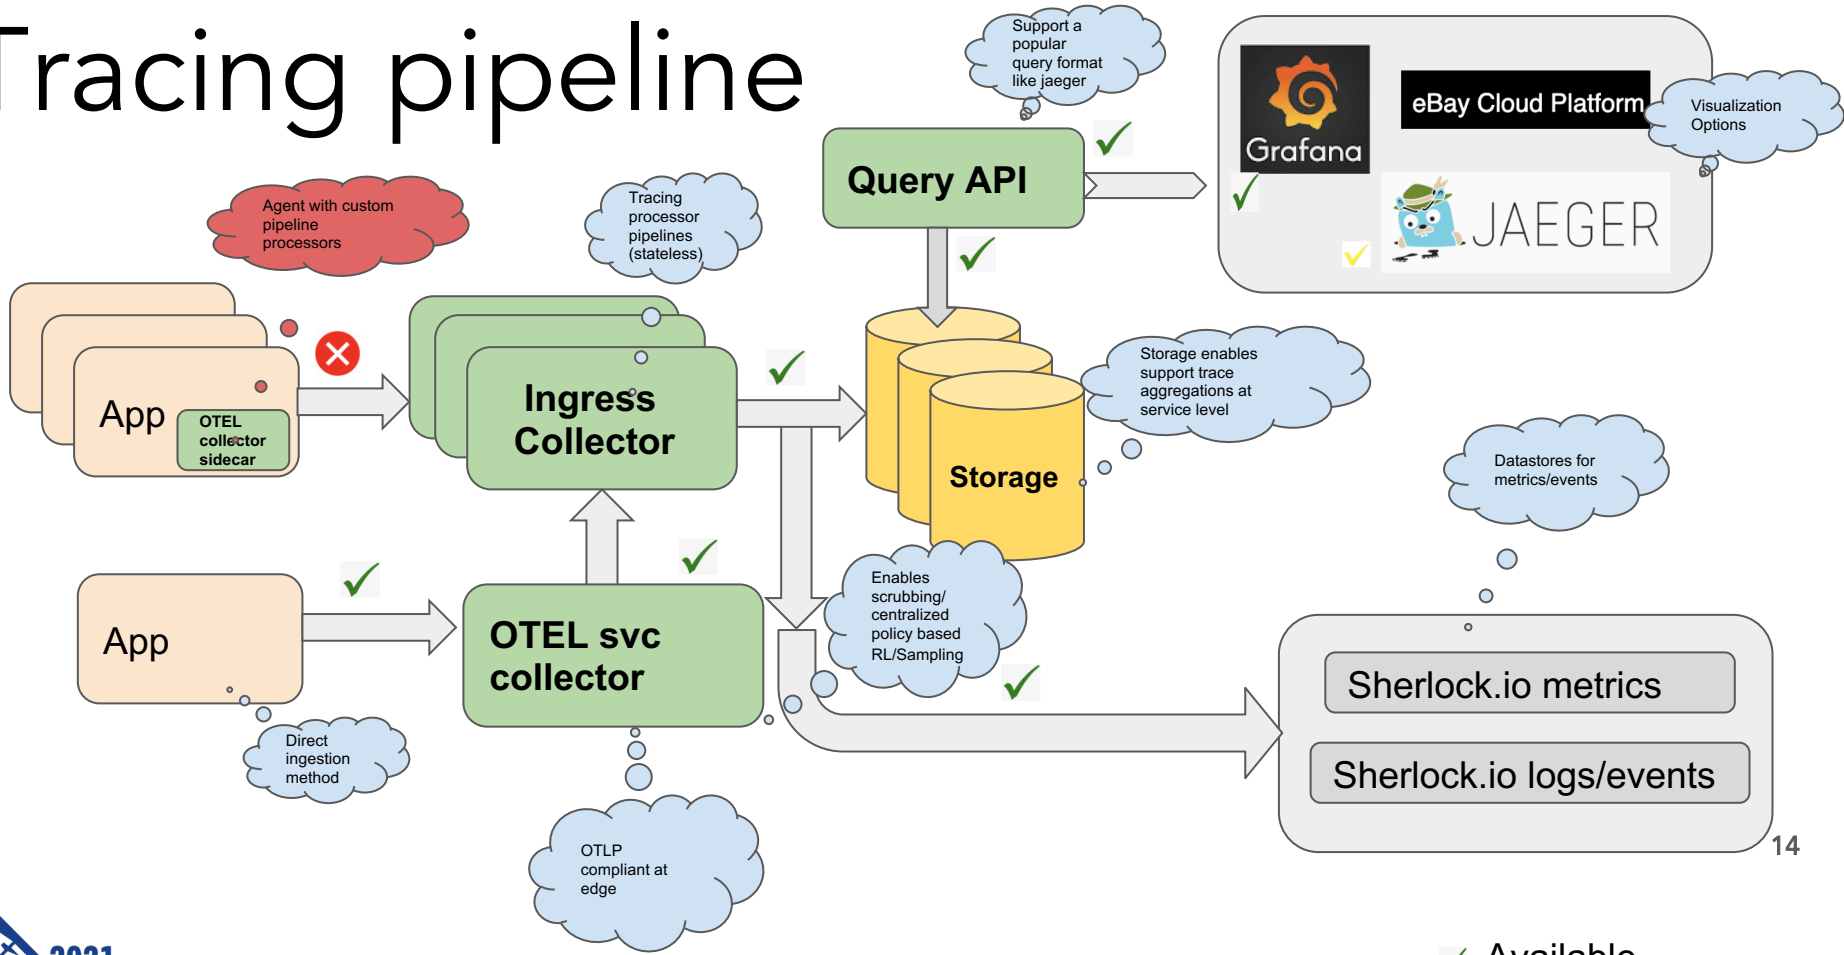

OSA CON 2021

✓ Available
✓ Upcoming

ClickHouse Schema

Tables	Purpose	Engine
traces	*Fetch traces for a given service & op*	ReplicatedMergeTree
tracessvc	*Given a Trace, find all associated services*	ReplicatedReplacingMergeTree
svcop	*services to op lookup*	ReplicatedReplacingMergeTree

```
CREATE TABLE sherlockio.traces
(
    `svc_name` LowCardinality(String),
    `operation` String,
    `traceid` String,
    `spanid` String,
    `pspanid` String,
    `start_time` DateTime,
    `status` LowCardinality(String),
    `httpstatus` LowCardinality(String),
    `log` Array(String) CODEC(ZSTD(1)),
    `duration` UInt64,
    `offset` UInt64,
    `attr_key` Array(String),
    `attr_value` Array(String),
    `attr_type` Array(UInt8),
    INDEX traceid traceid TYPE bloom_filter GRANULARITY 5,
    INDEX operation operation TYPE bloom_filter GRANULARITY 5
)
ENGINE = ReplicatedMergeTree('/clickhouse/{installation}/{cluster}/tables/{shard}/sherlockio/traces', '{replica}')
PARTITION BY toYYYYMMDDhhmmss(toStartOfHour(start_time))
ORDER BY (svc_name, traceid, start_time)
SETTINGS index_granularity = 8192
```

name	compressed	uncompressed
svc_name	24.05 MiB	4.99 GiB
status	438.64 MiB	89.52 GiB
start_time	4.41 GiB	19.90 GiB
operation	7.17 GiB	185.39 GiB
log	13.25 GiB	206.65 GiB
offset	22.18 GiB	39.81 GiB
traceid	23.19 GiB	164.20 GiB
attr_type	25.07 GiB	125.24 GiB
duration	25.34 GiB	39.81 GiB
pspanid	52.65 GiB	72.03 GiB
attr_key	66.01 GiB	1.46 TiB
spanid	83.97 GiB	84.59 GiB
attr_value	194.59 GiB	7.64 TiB

RED Metrics from Traces

Auto Instrumentation

- Requests
- Errors
- Duration
- HTTP status

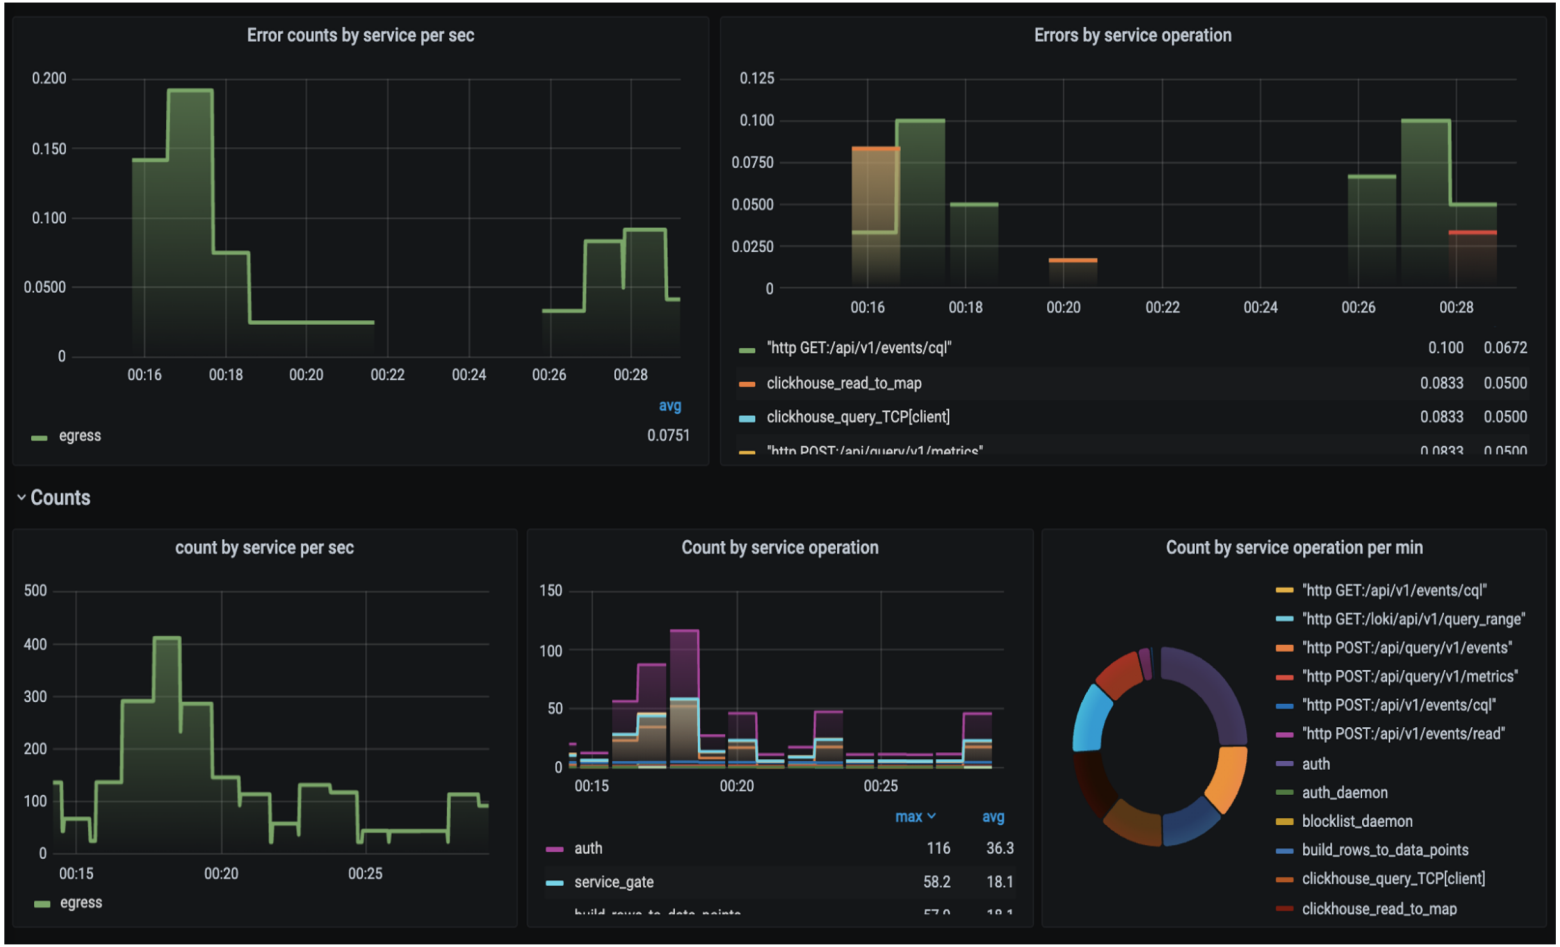

OSA CON 2021

Sampling

- Decision on whether to process/export a span or not
- Tail based sampling
- Sampling based on
  - *Latency*
  - *Error*
  - *HTTP response status codes*

Logs ⟺ Traces

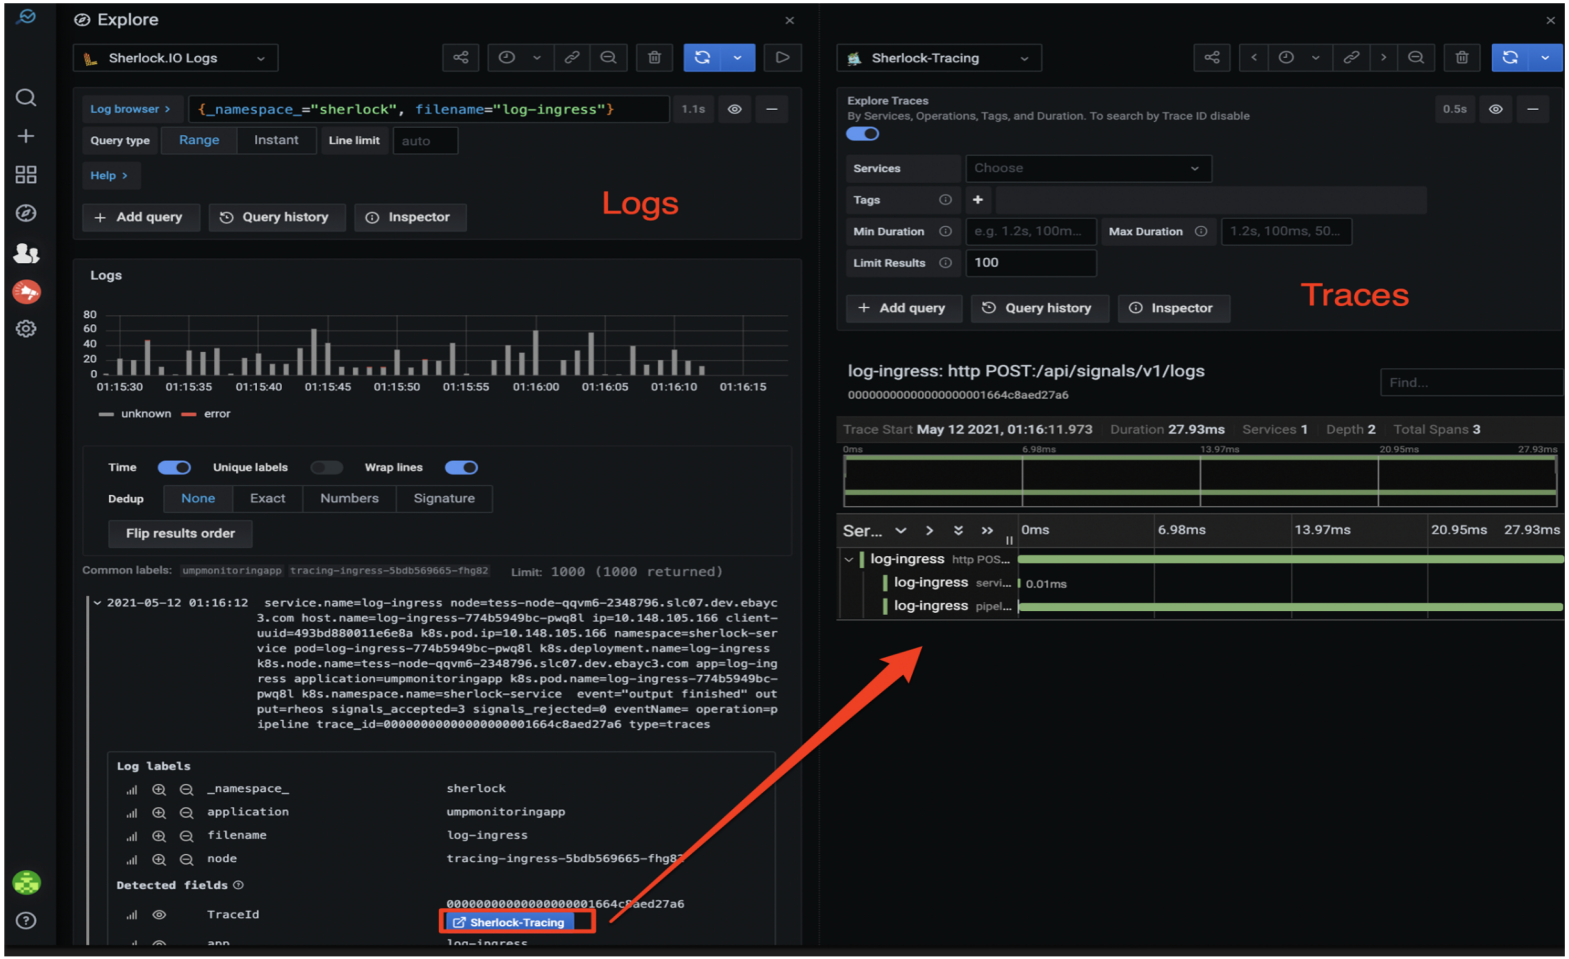
Explore
Sherlock.IO Logs
Log browser
{_namespace_="sherlock", filename="log-ingress"}
Query type
Range
Instant
Line limit
auto
Help
Add query
Query history
Inspector
Logs
Logs
80
60
40
20
0
01:15:30
01:15:35
01:15:40
01:15:45
01:15:50
01:15:55
01:16:00
01:16:05
01:16:10
01:16:15
unknown
error
Time
Unique labels
Wrap lines
Dedup
None
Exact
Numbers
Signature
Flip results order
Common labels:
Limit: 1000 (1000 returned)
2021-05-12 01:16:12 service.name=log-ingress node=tess-node-qqvm6-2348796.slc07.dev.ebayc3.com host.name=log-ingress-774b5949bc-pwq8l ip=10.148.105.166 client-uuid=493bd880011e6e8a k8s.pod.ip=10.148.105.166 namespace=sherlock-service pod=log-ingress-774b5949bc-pwq8l k8s.deployment.name=log-ingress k8s.node.name=tess-node-qqvm6-2348796.slc07.dev.ebayc3.com app=log-ingress application=umpmonitoringapp k8s.pod.name=log-ingress-774b5949bc-pwq8l k8s.namespace.name=sherlock-service event="output finished" output=rheos signals_accepted=3 signals_rejected=0 eventName= operation=pipeline trace_id=00000000000000000001664c8aed27a6 type=traces
Log labels
namespace sherlock
application umpmonitoringapp
filename log-ingress
node tracing-ingress-5bdb569665-fhg8
Detected fields
TraceId
00000000000000000001664c8aed27a6
Sherlock-Tracing
Sherlock-Tracing
Explore Traces
By Services, Operations, Tags, and Duration. To search by Trace ID disable
Services
Choose
Tags
Min Duration
e.g. 1.2s, 100m...
Max Duration
1.2s, 100ms, 50...
Limit Results
100
Add query
Query history
Inspector
Traces
log-ingress: http POST:/api/signals/v1/logs
00000000000000000001664c8aed27a6
Find...
Trace Start May 12 2021, 01:16:11.973
Duration 27.93ms
Services 1
Depth 2
Total Spans 3
0ms
6.98ms
13.97ms
20.95ms
27.93ms
Ser...
log-ingress http POS...
log-ingress servi... 0.01ms
log-ingress pipel...

Metrics ⇒ Traces

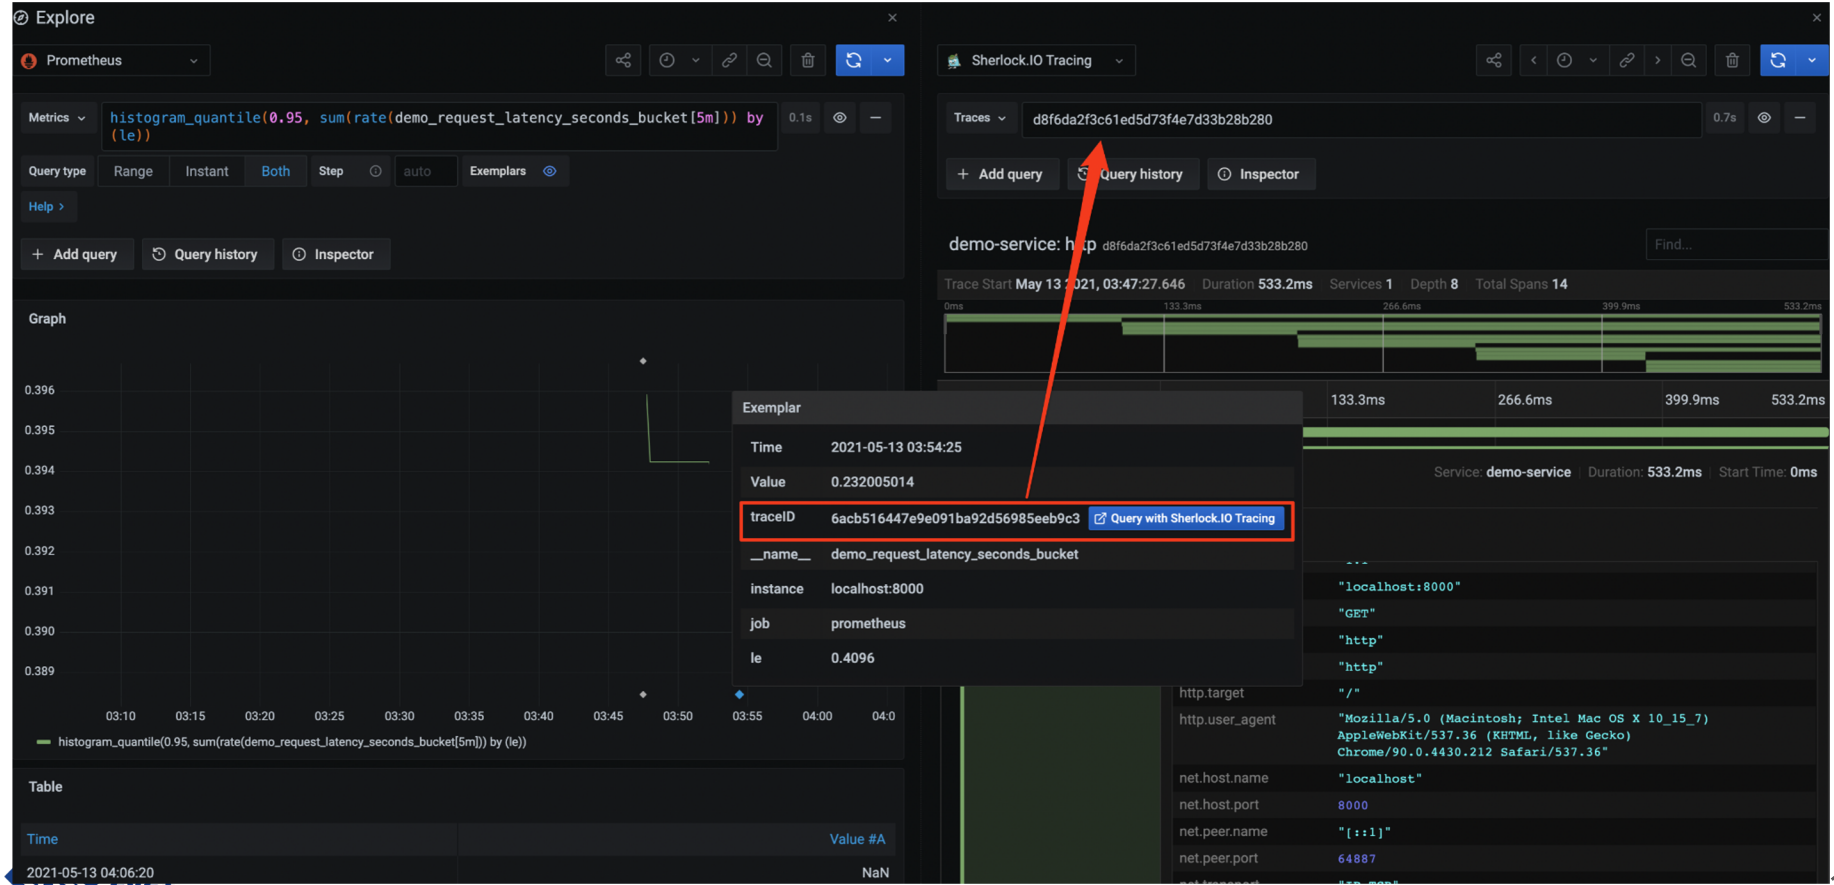

Challenges

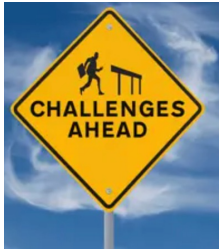

- Adoption
  - *Frmk integration*
- Coverage
  - *Service Mesh*
- Sampling

eBay is Hiring

- Globally across all major cities and countries.
- Remote Work is an potential option in current times.
- Areas
  - RCGs
  - Research and Engineering
  - Data Science
  - Product Management
- Explore opportunities at https://careers.ebayinc.com
- Reach out to amvaidya@ebay.com or sudekumar@ebay.com

Thank You!