

HACKING OPENTELEMETRY

Alberto Gastaldello, Andrea Caretta - Liquid Reply IT

Conf42, June 8th 2023


AGENDA

- 1 Intro – Observability tips
- 2 OpenTelemetry
- 3 Distributed Traces
- 4 Hacking OpenTelemetry
- 5 Demo

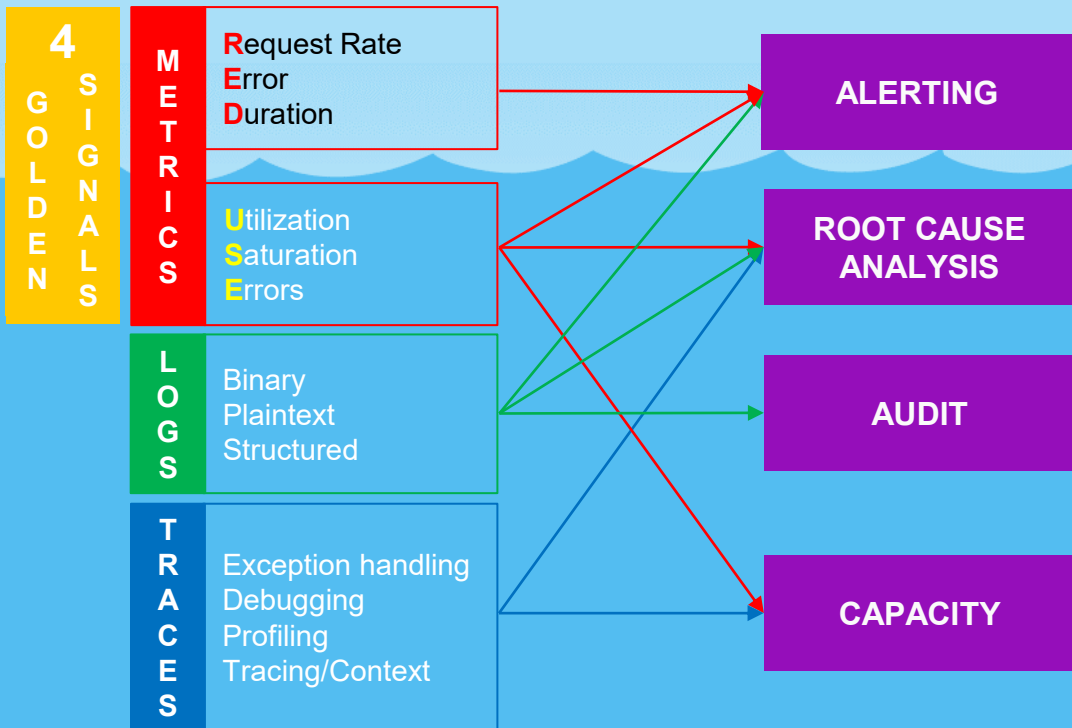


OBSERVABILITY PILLARS

Symptoms



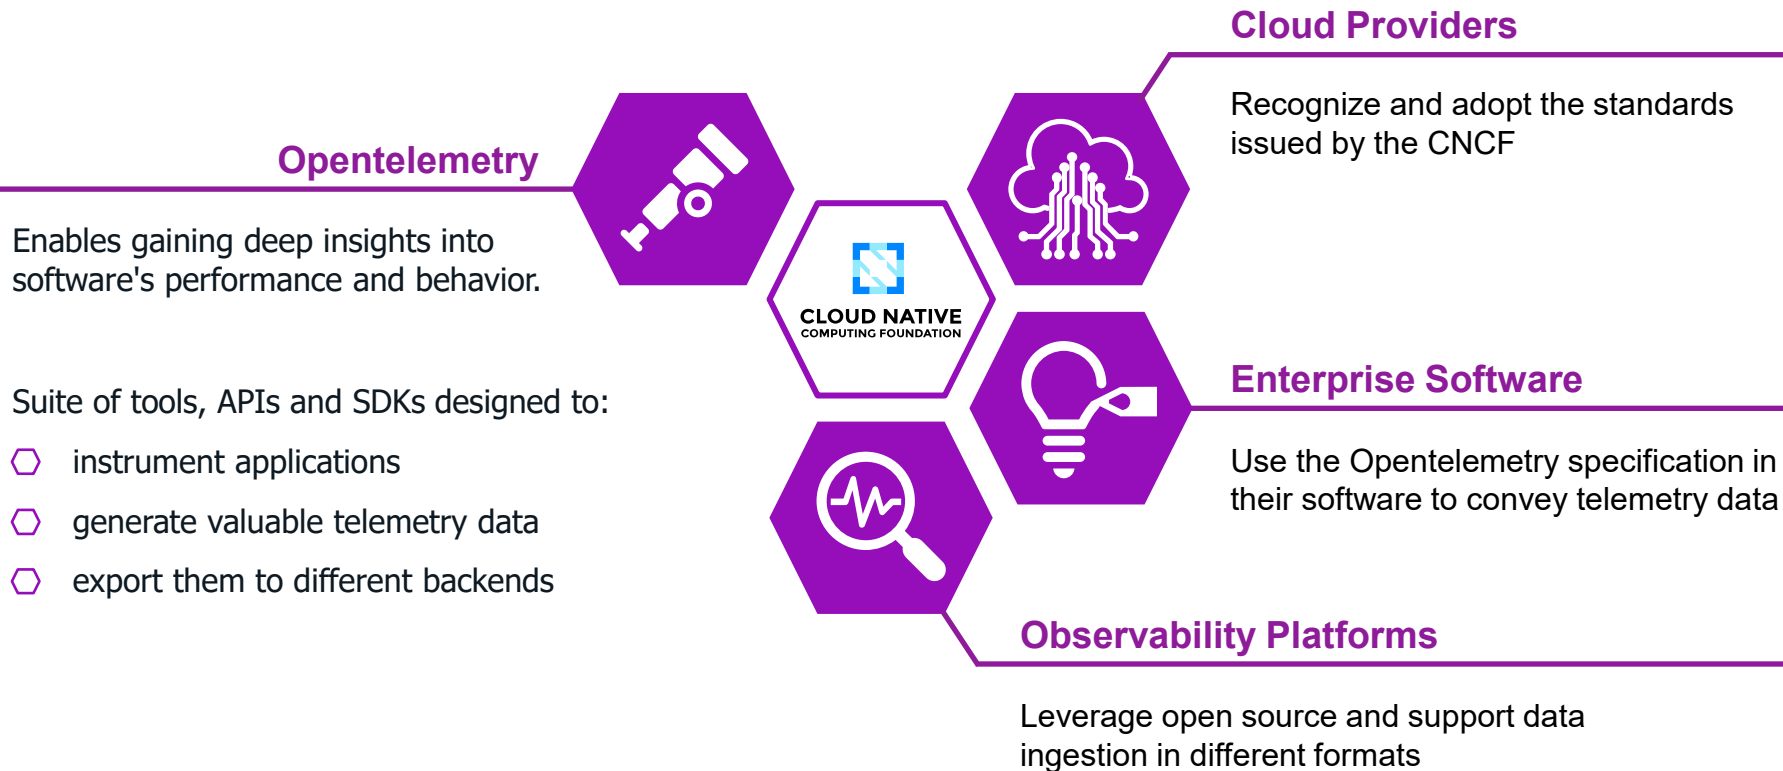
Causes



OPEN TELEMETRY



ADOPTERS AND CONTRIBUTORS



COMPONENTS

Language-specific SDK

They allow you to use APIs to generate telemetry data, export it to different backends propagating context information

Auto-instrumentation agents

Packages capable of generating telemetry data without changing the application code; many libraries are supported and recognized automatically

K8s Operator


Kubernetes Operator implementation, manages Otel Collector and self-instrumentation of workloads in a distributed environment

OpenTelemetry Collector

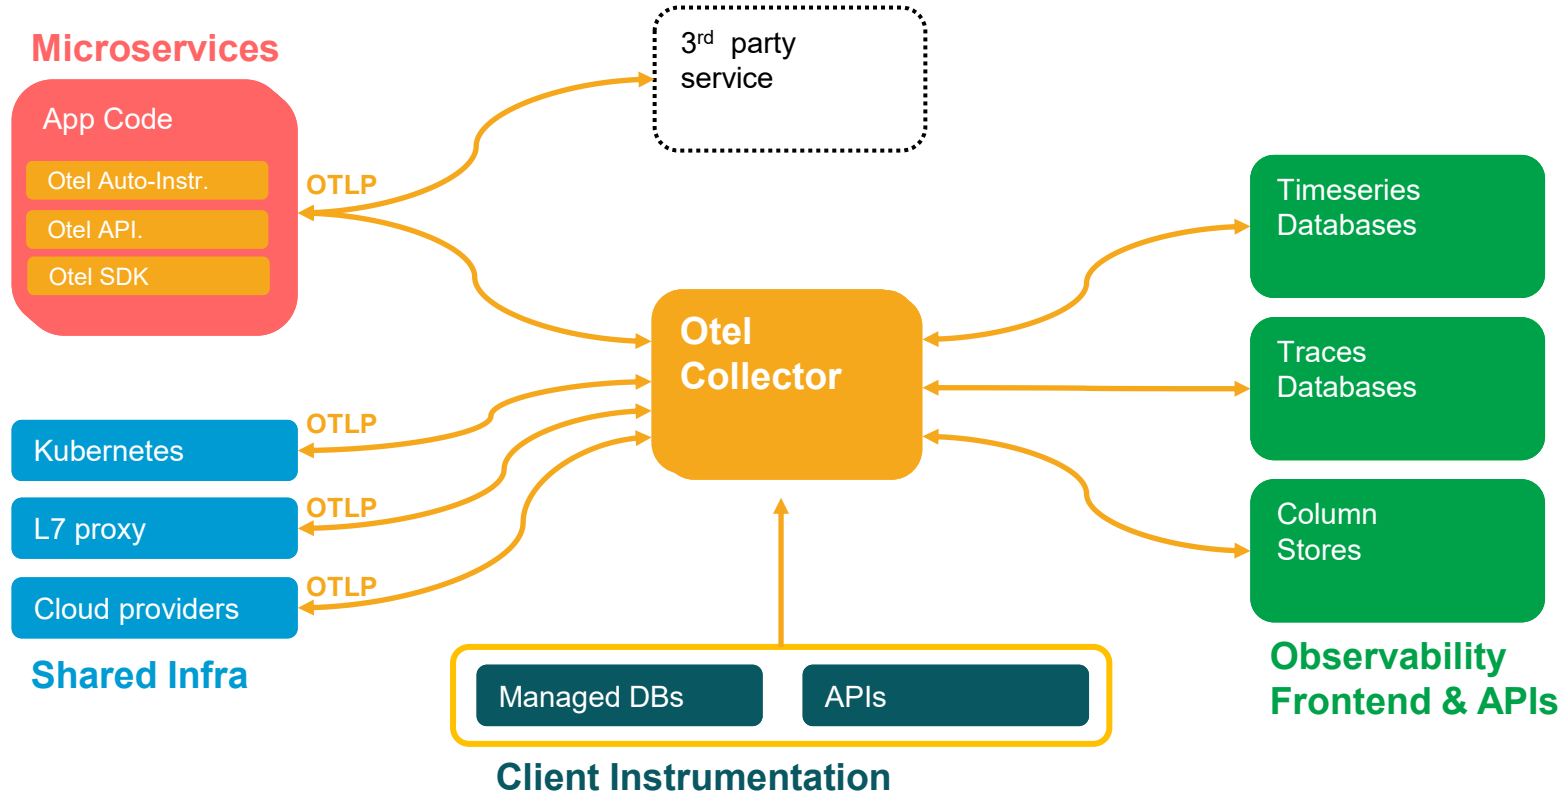
Proxy capable of receiving telemetry data in different formats, editing and/or filtering it, and sending it in different formats to the desired backend

Cross-language references

API (data types and operations),
Data (semantic conventions)



DATA FLOW



DISTRIBUTED TRACES



TRACING CONCEPTS

TRACE

A "distributed trace", or simply "trace", **records the paths followed by a request** (triggered by an application or an end-user) propagating through multi-service architectures.

A trace consists of one or more **spans**.

SPAN

A **span** is identified as a trace unit.

As a building block of a distributed trace, it **traces specific operations that a request makes**, reporting a drawing of what happened during the interval in which the operation was performed.

CONTEXT

Defines a universally recognized format for **propagating data and linking them together**. Can be adopted by different platforms and different tracking tools.

The following headers are propagated:

- **traceparent**
- **tracestate**



TRACING CONCEPTS

TRACE

SPAN

CONTEXT

traceparent = 00-0af7651916cd43dd8448eb211c80319c-b9c7c989f97918e1-01

TRACE ID

SPAN ID

PARENT SPAN ID

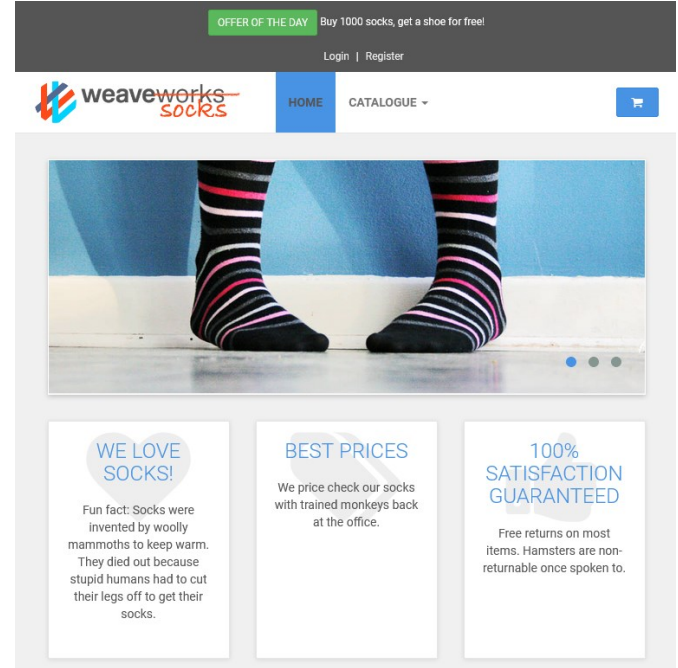
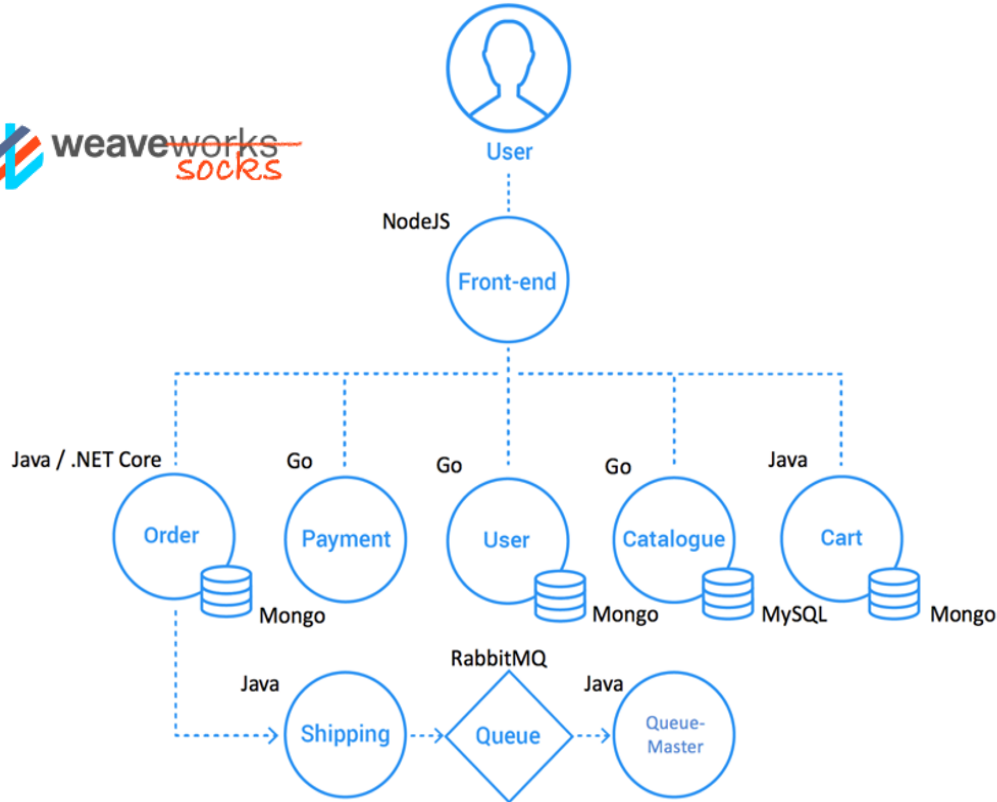
Identifier of the whole trace tree

Identifier of the current span

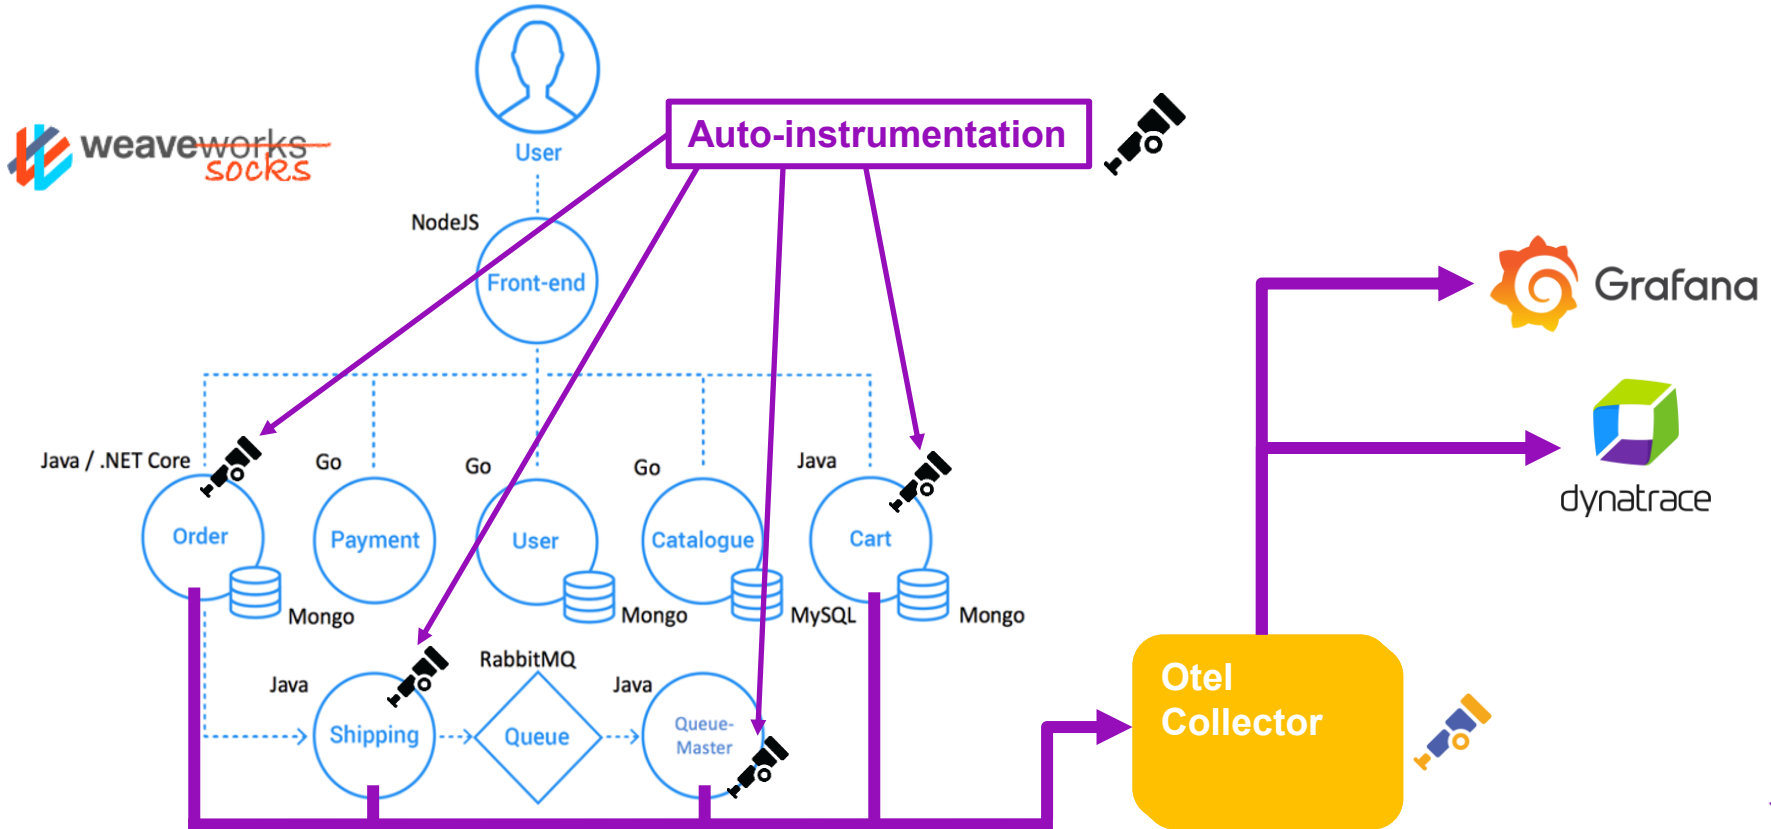
Identifier of the span
generated by the caller



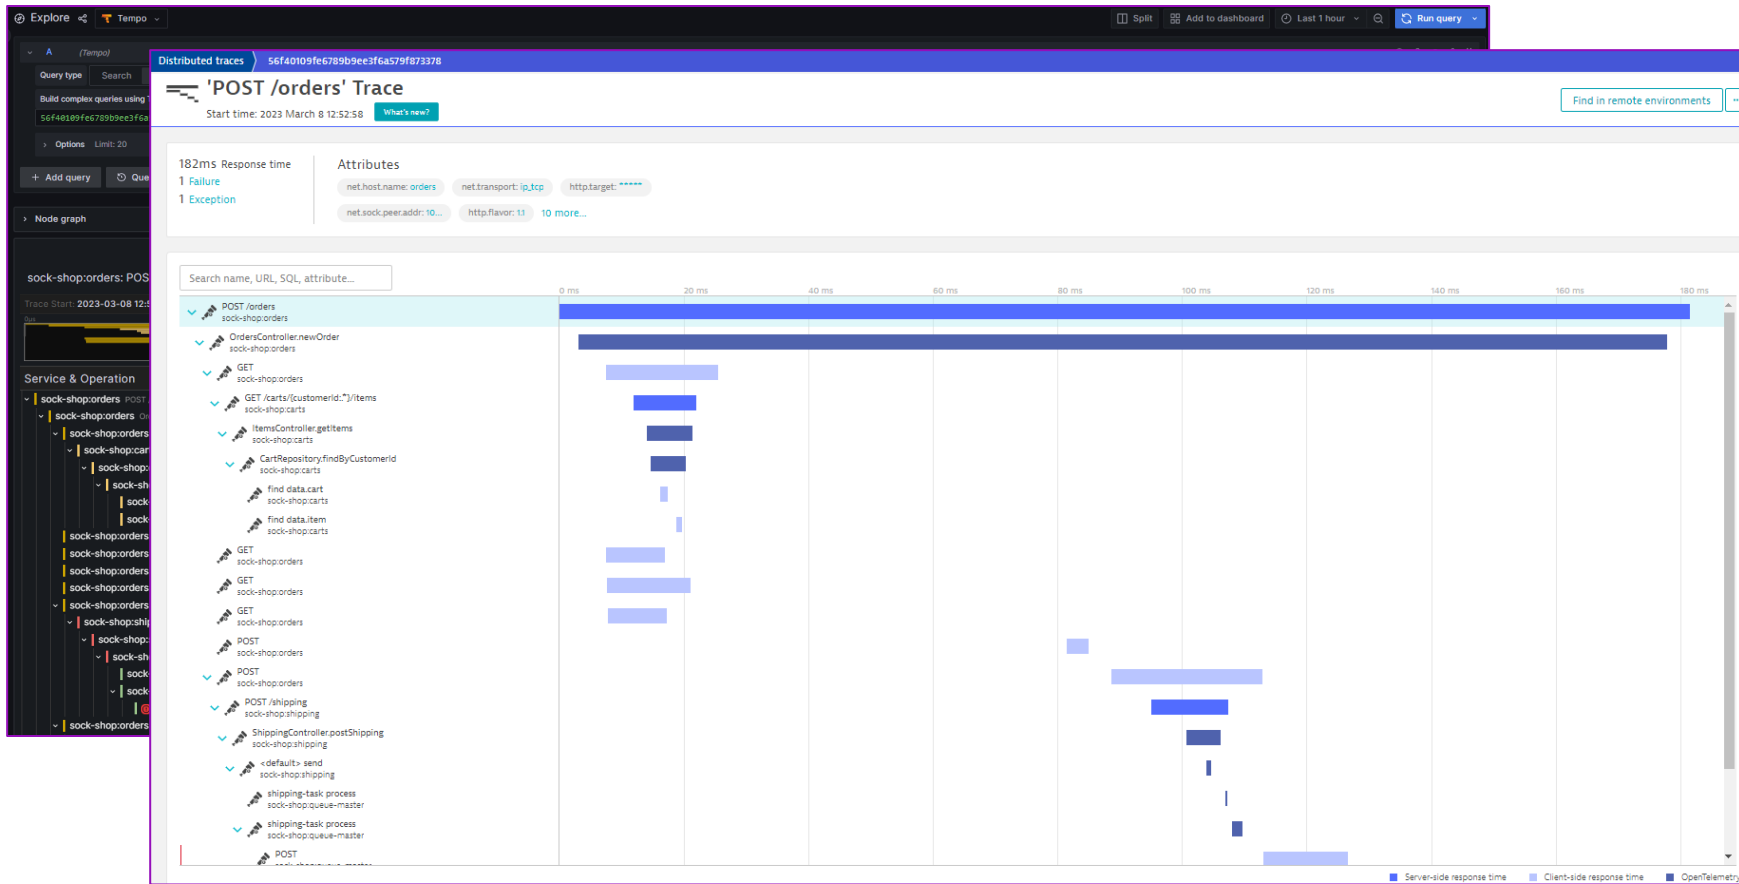
A COMMUNICATING SCENARIO



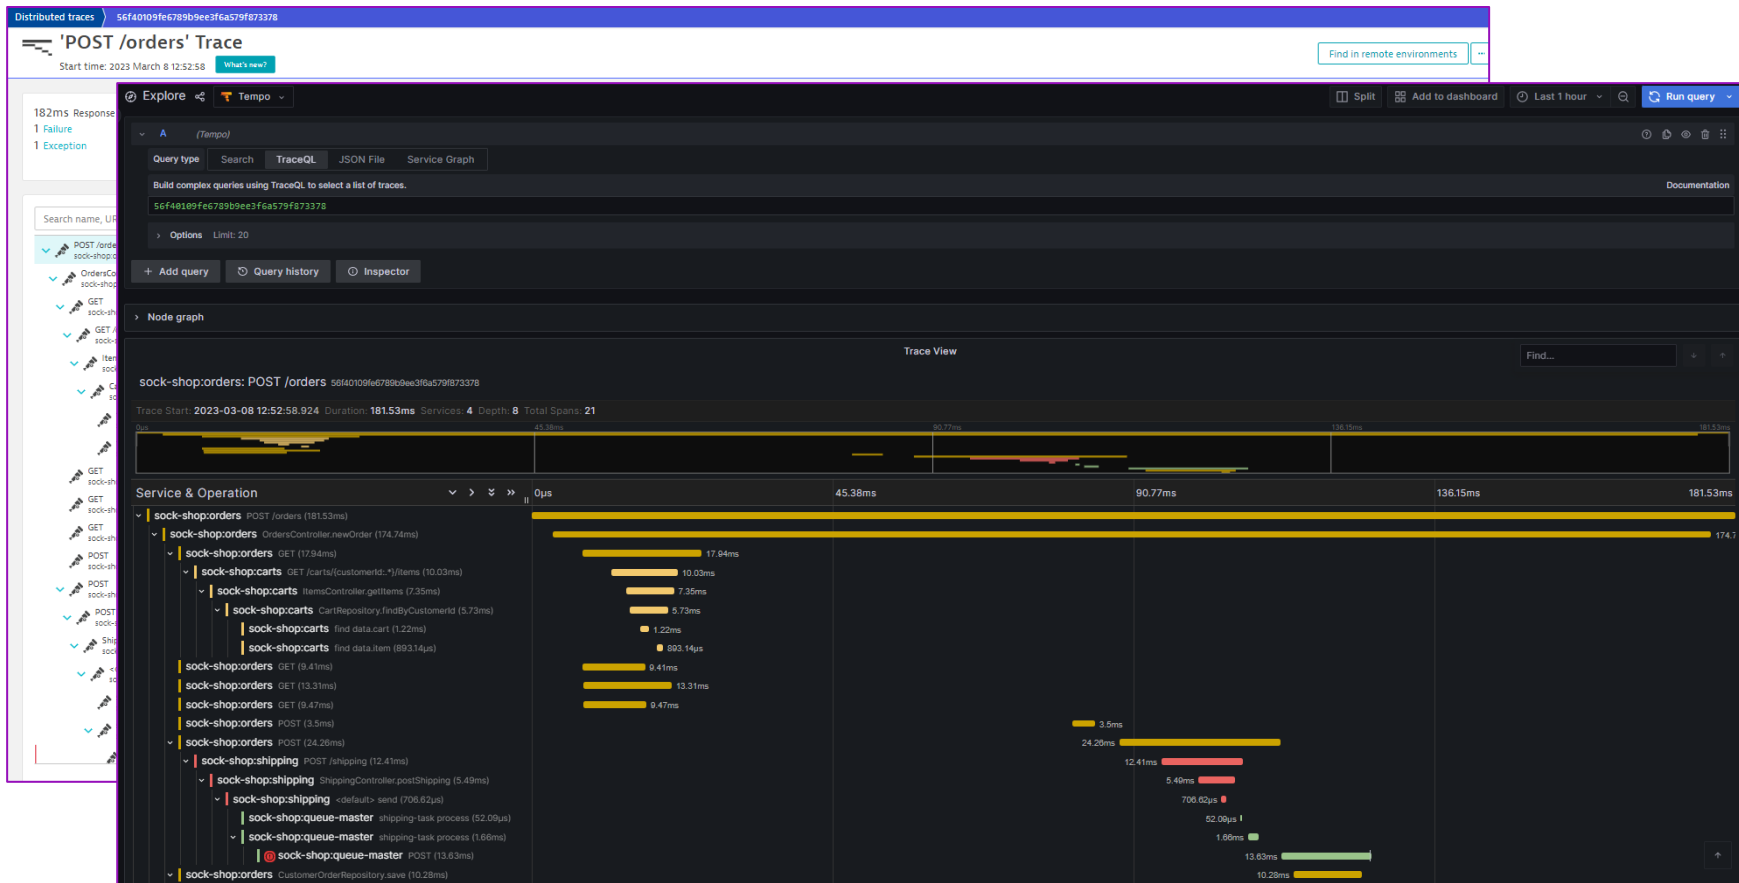
FOLLOWING A TRANSACTION



TRACKING RESULTS - DYNATRACE



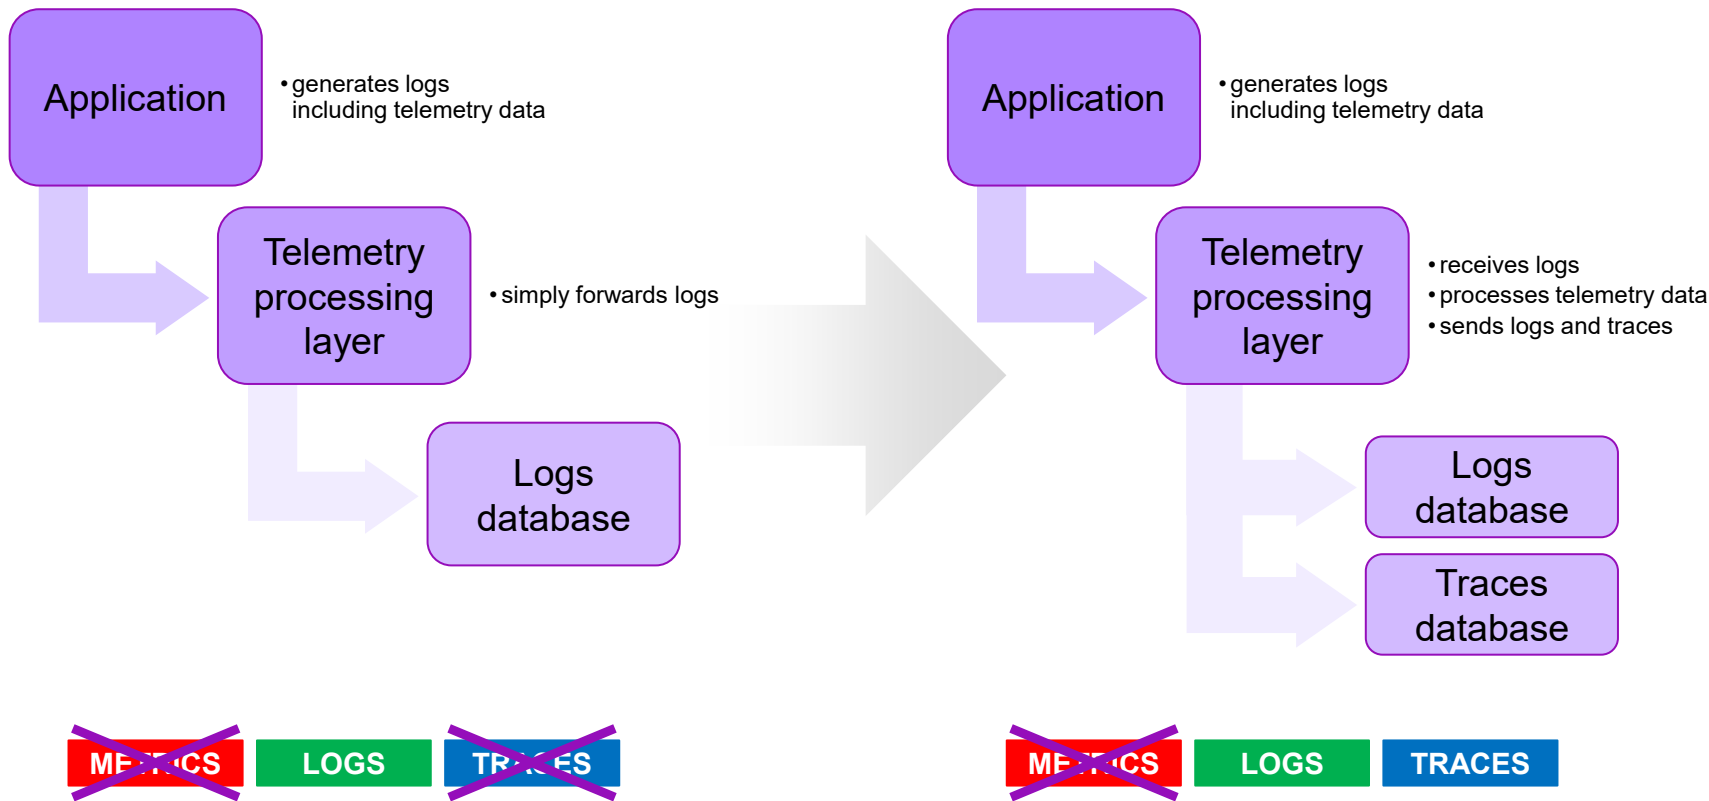
TRACKING RESULTS - GRAFANA



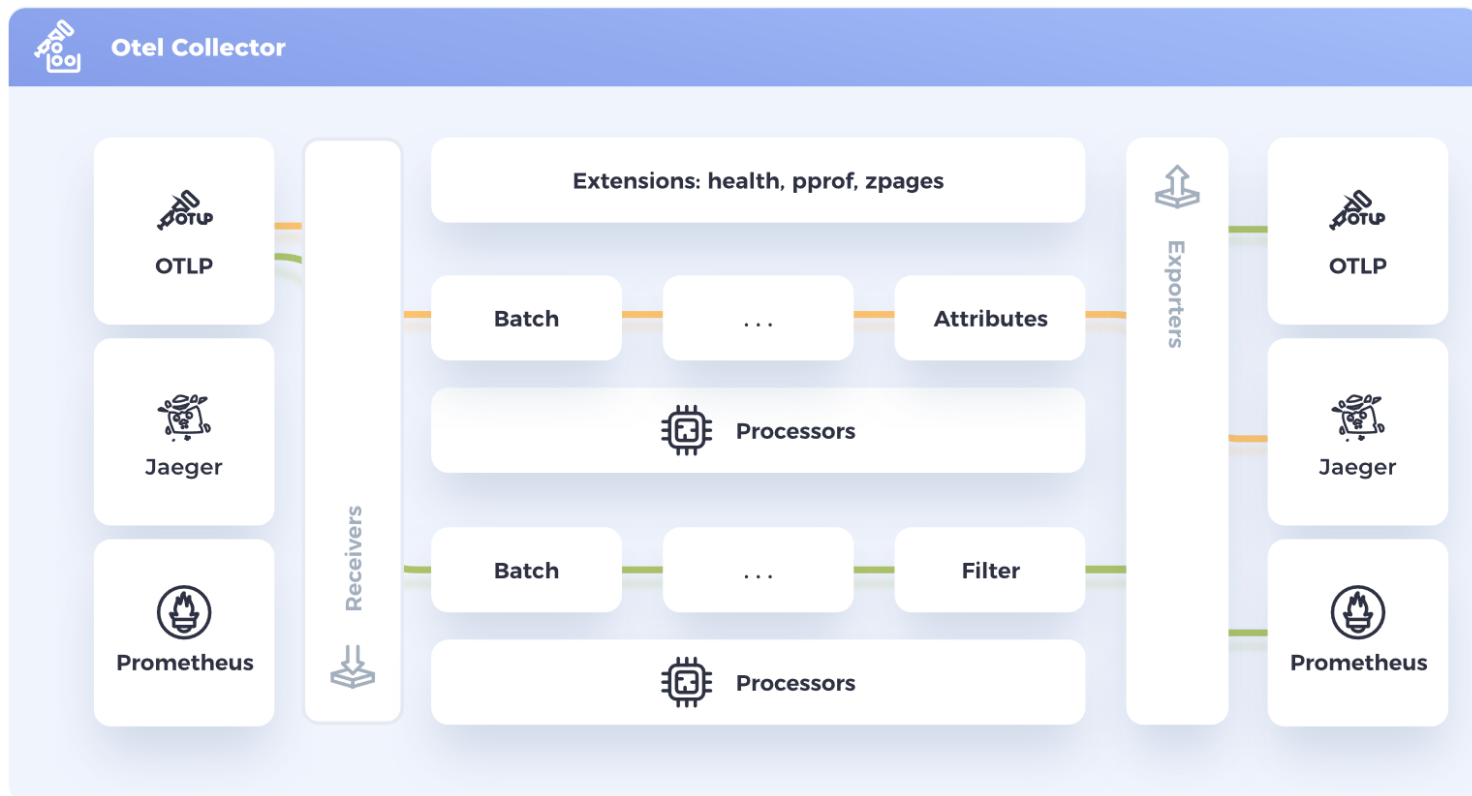
HACKING THE COLLECTOR



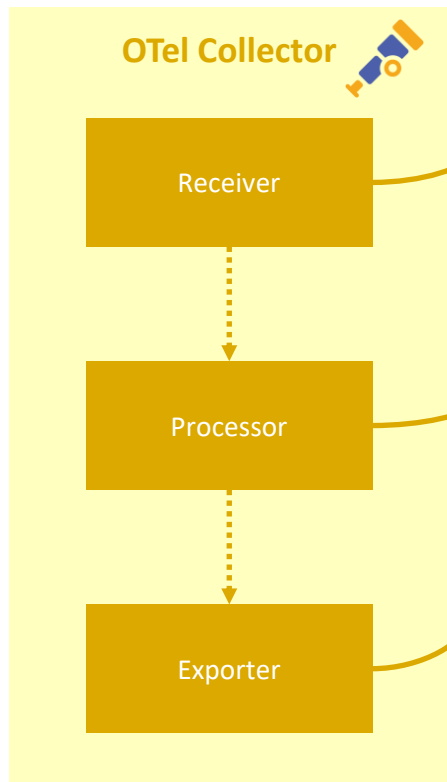
ONLY LOGS AVAILABLE?



COLLECTOR MODULES & PIPELINES



COLLECTOR CONFIGURATION



```
receivers:  
  otlp:  
    protocols:  
      grpc:  
      http:  
    syslog:  
    udp:  
      listen_address: "0.0.0.0:54526"  
    protocol: rfc5424
```

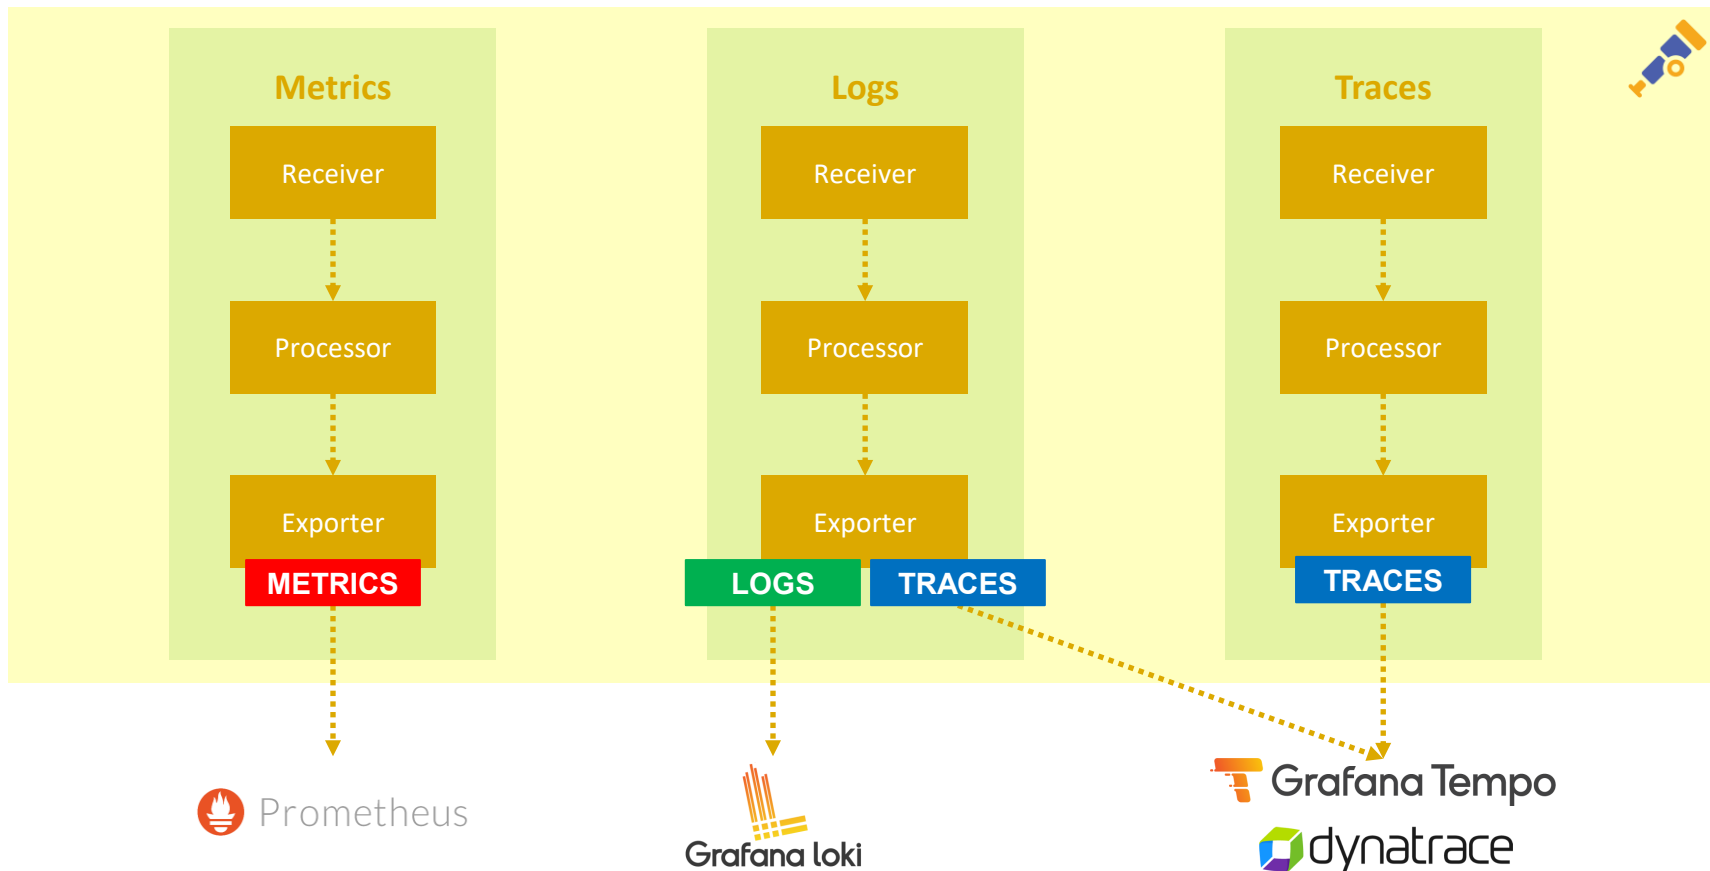
```
processors:  
  transform:  
    log_statements:  
      - context: log  
      statements:  
        - set(severity_text, "FAIL") where body == "request failed"  
        - replace_all_matches(attributes, "/user/*/list/*", "/user/{userId}/list/{listId}")  
        - replace_all_patterns(attributes, "/account/\\d{4}", "/account/{accountId}")  
        - set(body, attributes["http.route"])
```

```
exporters:  
  loki:  
    endpoint: "https://user:token@logs-prod-example.grafana.net/loki/api/v1/push"
```

```
service:  
  pipelines:  
    logs:  
      receivers: [otlp, syslog]  
      processors: [transform]  
      exporters: [loki]
```



FROM LOG TO TRACE



EVERYTHING CAN BE A SPAN

Application

Telemetry
processing
layer

LOG

```
<34>1 2022-12-16T11:15:57.000Z mymachine.example.com su - ID47  
[opentelemetry trace_id="0123456789abcdef0123456789abcdef"  
span_id="0123456789abcdef" parent_span_id="abcdef0123456789"  
trace_flags="01"] BOM'su root' failed for lonvick on /dev/pts/8
```

LOG

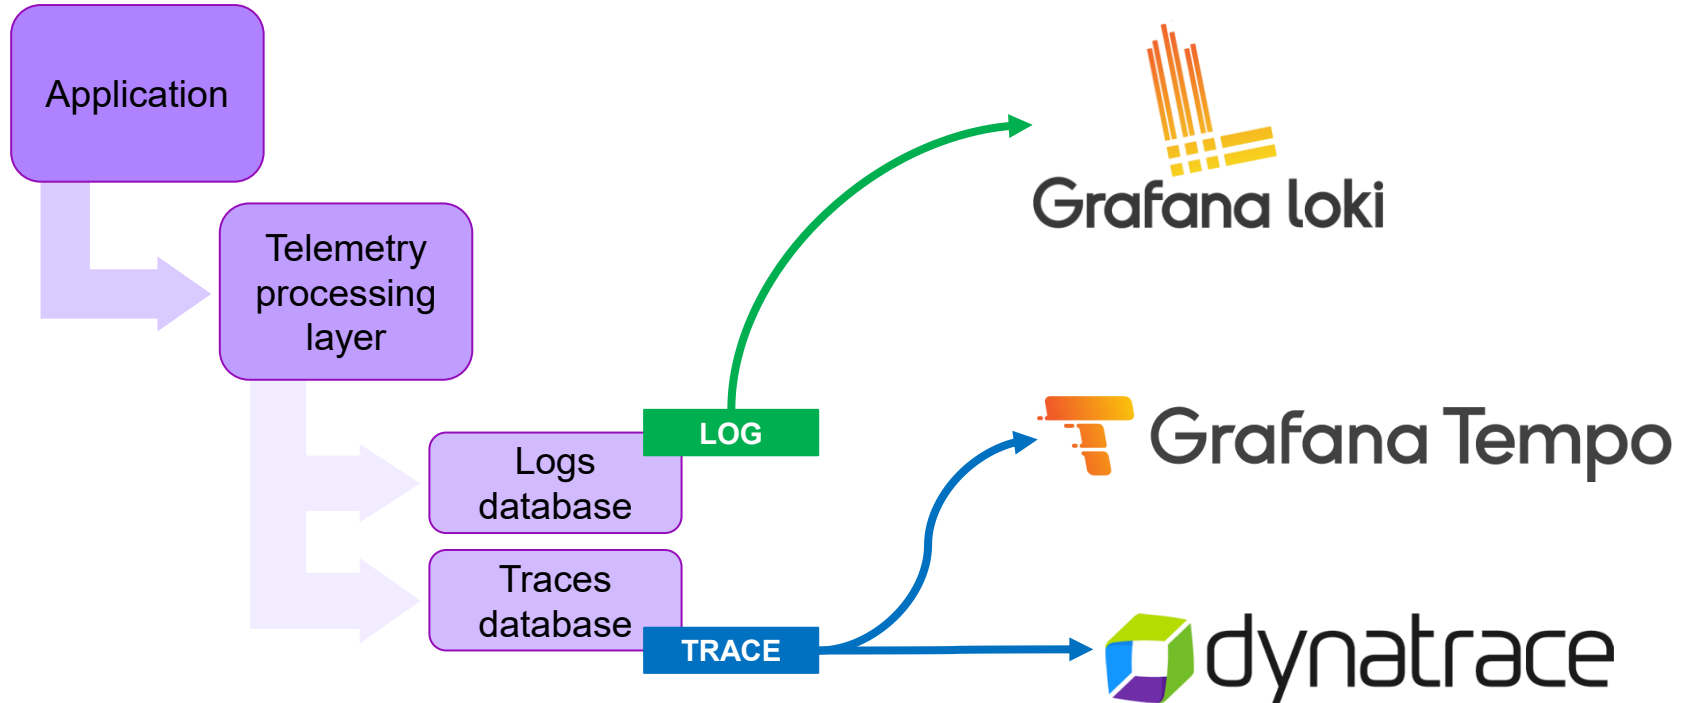
```
{  
  "timeUnixNano": "1671189357000000000",  
  "observedTimeUnixNano": "1671189357107662600",  
  "severityNumber": 18,  
  "severityText": "crit",  
  "body": {"stringValue": "<34>1 2022-12-16T11:15:57.000Z  
mymachine.example.com su - ID47 [opentelemetry  
trace_id=\\"0123456789abcdef0123456789abcdef\  
span_id=\\"0123456789abcdef\" parent_span_id=\\"abcdef0123456789\  
trace_flags=\\"01\\\" BOM'su root' failed for lonvick on /dev/pts/8\"),  
  "attributes": [  
    {  
      "key": "parentSpanId",  
      "value": {"stringValue": "abcdef0123456789"}  
    },  
    ...  
  ],  
  "traceId": "0123456789abcdef0123456789abcdef",  
  "spanId": "0123456789abcdef"  
}
```

TRACE SPAN

```
{  
  "trace_id": "0123456789abcdef0123456789abcdef",  
  "span_id": "0123456789abcdef",  
  "parent_span_id": "abcdef0123456789",  
  "start_time_unix_nano": 1671189357000000000,  
  "end_time_unix_nano": 1671189357107662600,  
  "attributes": [  
    {  
      "key": "body",  
      "value": {  
        "string_value": "<34>1 2022-12-16T11:15:57.000Z  
mymachine.example.com su - ID47 [opentelemetry  
trace_id=\\"0123456789abcdef0123456789abcdef\  
span_id=\\"0123456789abcdef\" parent_span_id=\\"abcdef0123456789\  
trace_flags=\\"01\\\" BOM'su root' failed for lonvick on /dev/pts/8"  
      }  
    }  
  ]  
}
```



TELEMETRY DATA STORAGE



DEMO



CONTACTS

Alberto Gastaldello

al.gastaldello@reply.it

+39 342 864 7689

Andrea Caretta

a.caretta@reply.it

+39 345 970 7769



REPLY
LIQUID

Liquid Reply IT
IT Services and IT Consulting · Turin, Piedmont · 704 followers

Gianluciano & 36 other connections work here · 99 employees

✓ Following Invite More



