



**ENTERPRISE
DEVOPS**
T E C H C O N

Debugging Microservices in Kubernetes

Sander Rodenhuis




Debugging Microservices in Kubernetes

Using Istio, OpenTelemetry and Grafana Tempo

Agenda

- The 3 pillars of observability
- The what and why of Distributed tracing
- Using Istio for tracing (the good, the .. and the ...)
- Introducing OpenTelemetry
- Things to consider when getting started
- Observability with distributed tracing in Otomi
- Demo: Tracing with Nginx, Istio, OpenTelemetry, Prometheus, Grafana Loki and Tempo

The 3 pillars of Observability



The 3 pillars of Observability

Logs

Records of events, warnings and errors within a microservice.

Provide insights into events and errors during the lifecycle of microservices.

Metrics

Quantifiable measurements that reflect the health and performance of a microservice.

Provide real-time insight into the state of microservices.

Traces

Data that tracks a request as it flows through various microservices.

Research the root cause of a problem with a microservices architecture consisting out of multiple microservices.

Logs, Metrics and Traces each provide valuable but limited visibility. Only by combining them you'll get the complete picture. And everything is about CONTEXT!

Distributed tracing

- A method of tracking application requests as they flow between microservices
- Understand the behaviour of an application consisting out of multiple microservices
- Troubleshooting performance bottlenecks
- Fix errors, and other issues that could impact the user experience
- But you can't trace everything!

The basics

Trace

A complete **end-to-end path of a request** or transaction as it flows through a distributed system.

Represents the journey of a specific operation as it traverses various components and services in a distributed architecture.

Span

A single operation or **unit of work** within a distributed system.

Captures the timing and metadata associated with a specific operation and provides a way to track and understand the behavior of individual components and services.

Context Propagation

Passing contextual information between different components or services within a distributed system.

Crucial for connecting and correlating spans to construct a complete trace of a request or transaction as it flows through various services.

Why tracing with Istio

- Responsible for managing traffic, it can also report logs, metrics, and traces
- Leverages Envoy's distributed tracing feature to provide tracing integration out of the box
- A service mesh can introduce its own delays and issues
- Visibility into this layer of infrastructure is useful in troubleshooting
- Easy to get started

Tracing using Istio (defaults)

The good

- Easy to setup
- No coding required

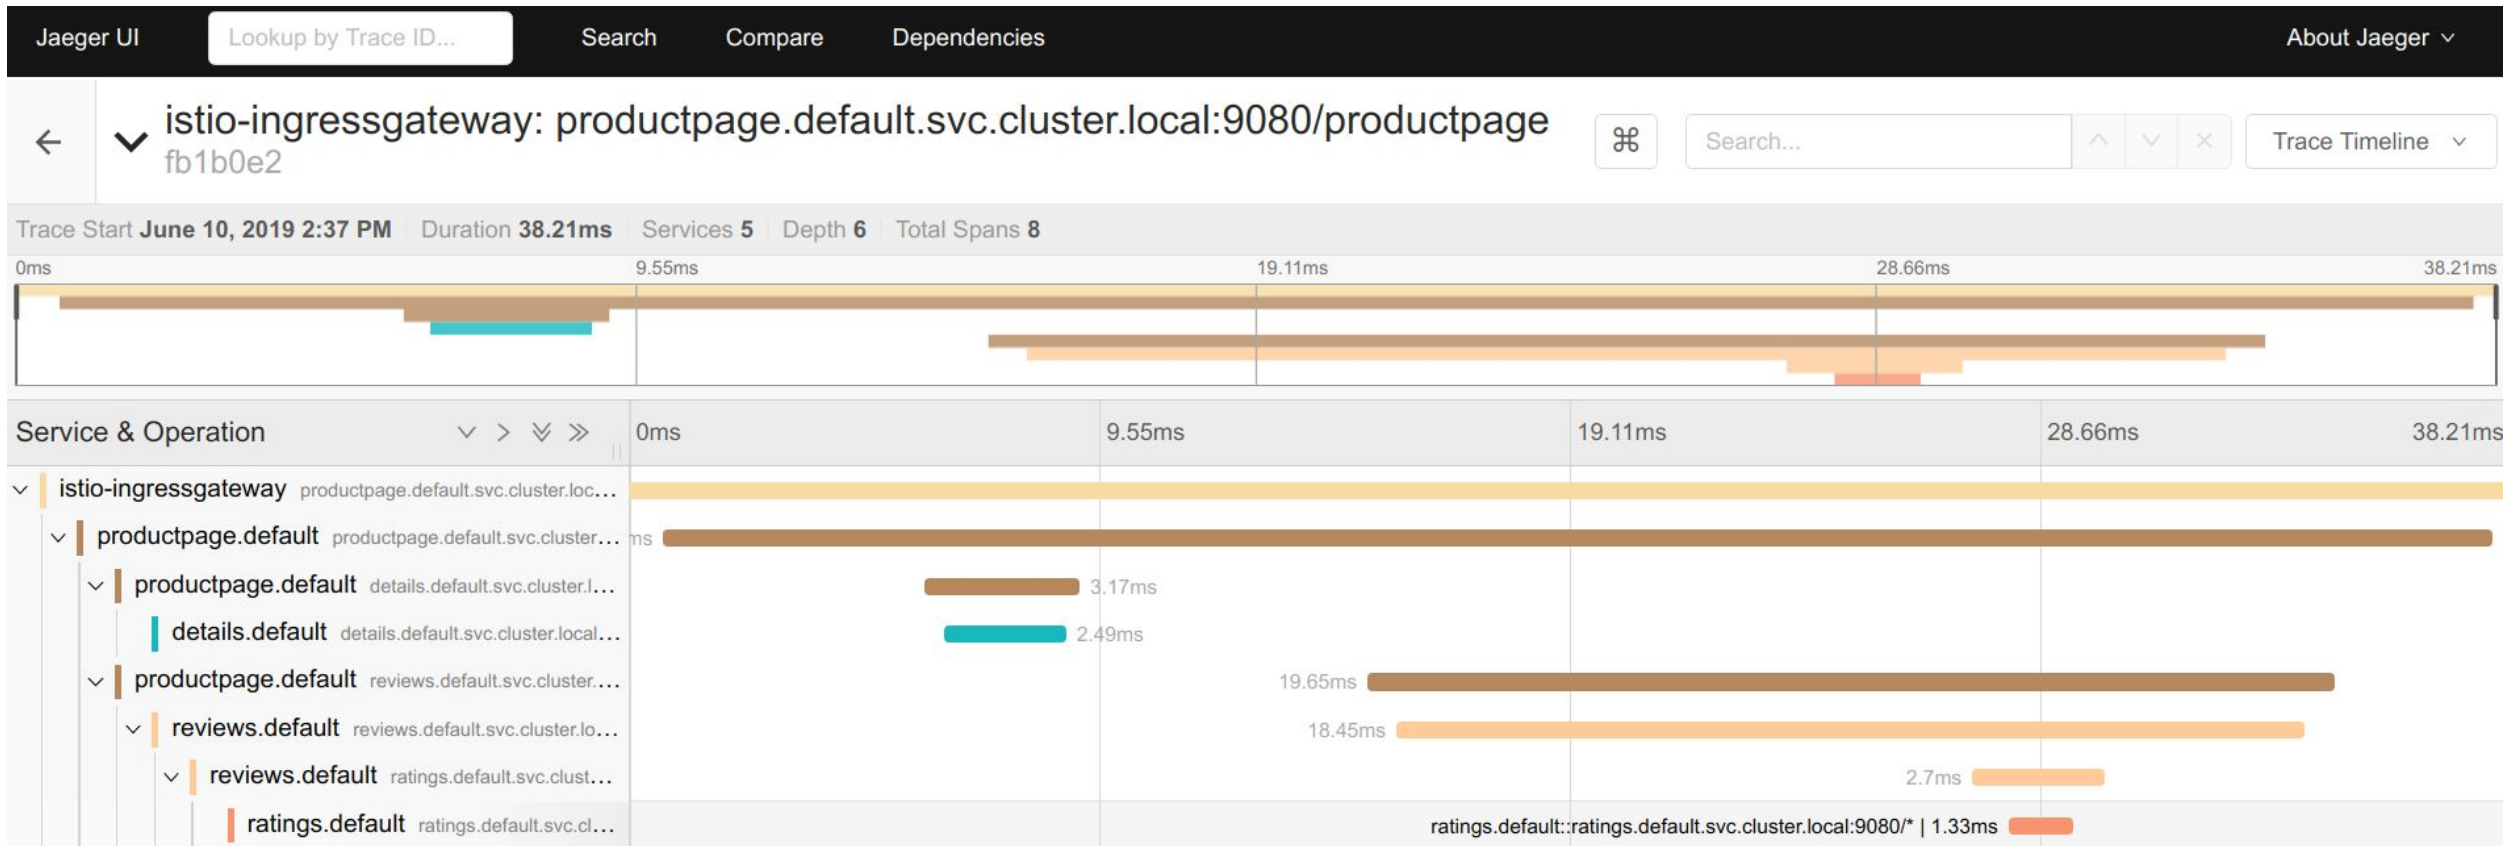
The bad

- Only collects partial data with partial context
- One sampling rate for all traffic

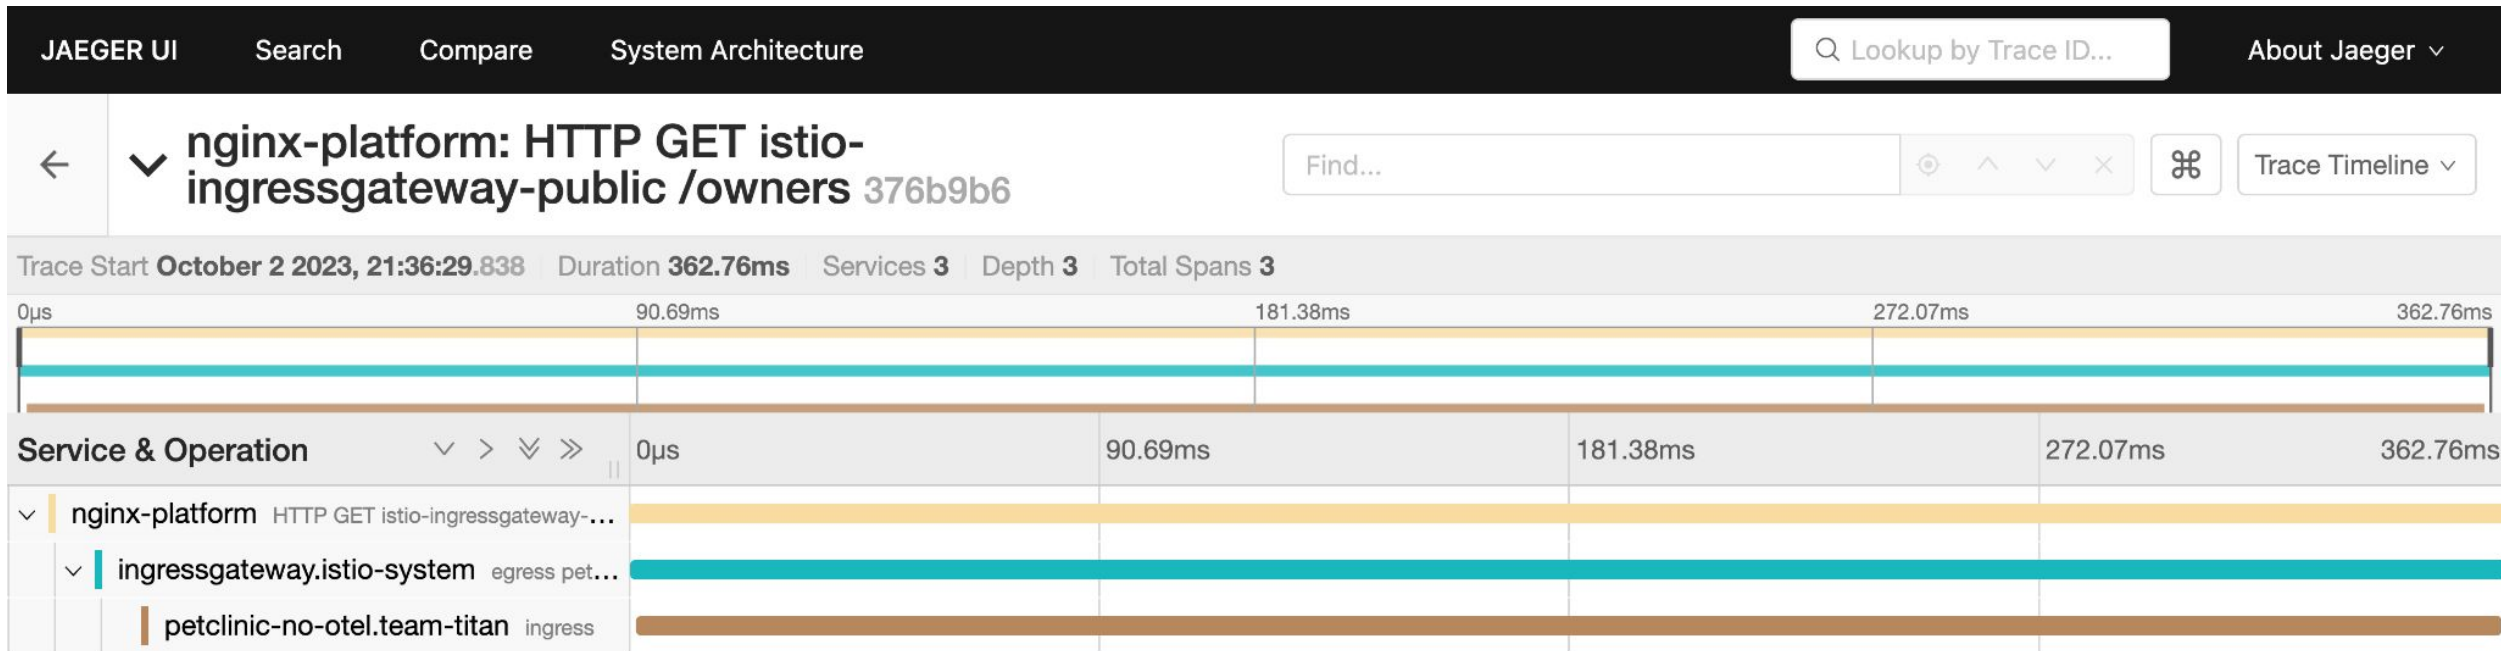
The ugly

- Uses Jaeger and Zipkin format
- No active development for Jaeger SDK (needed for instrumentation)

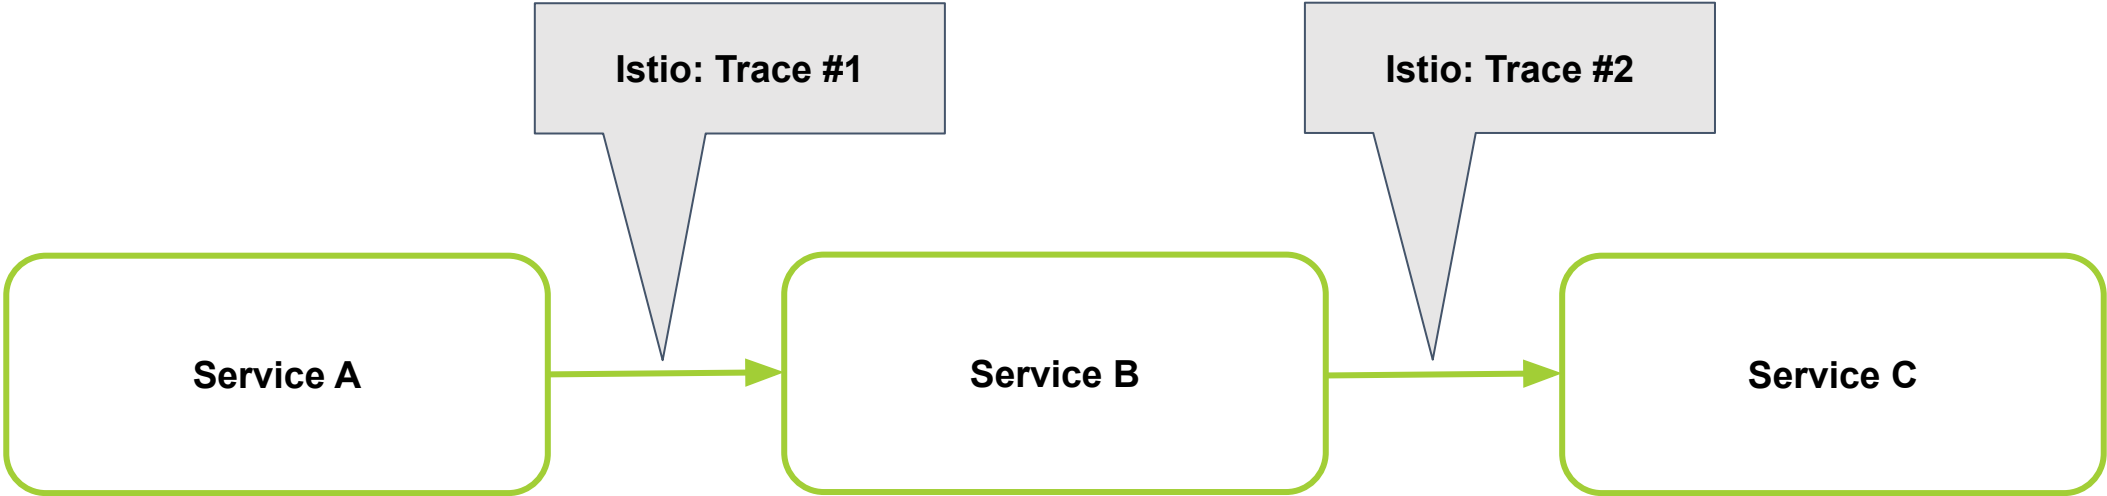
What we see in the demo



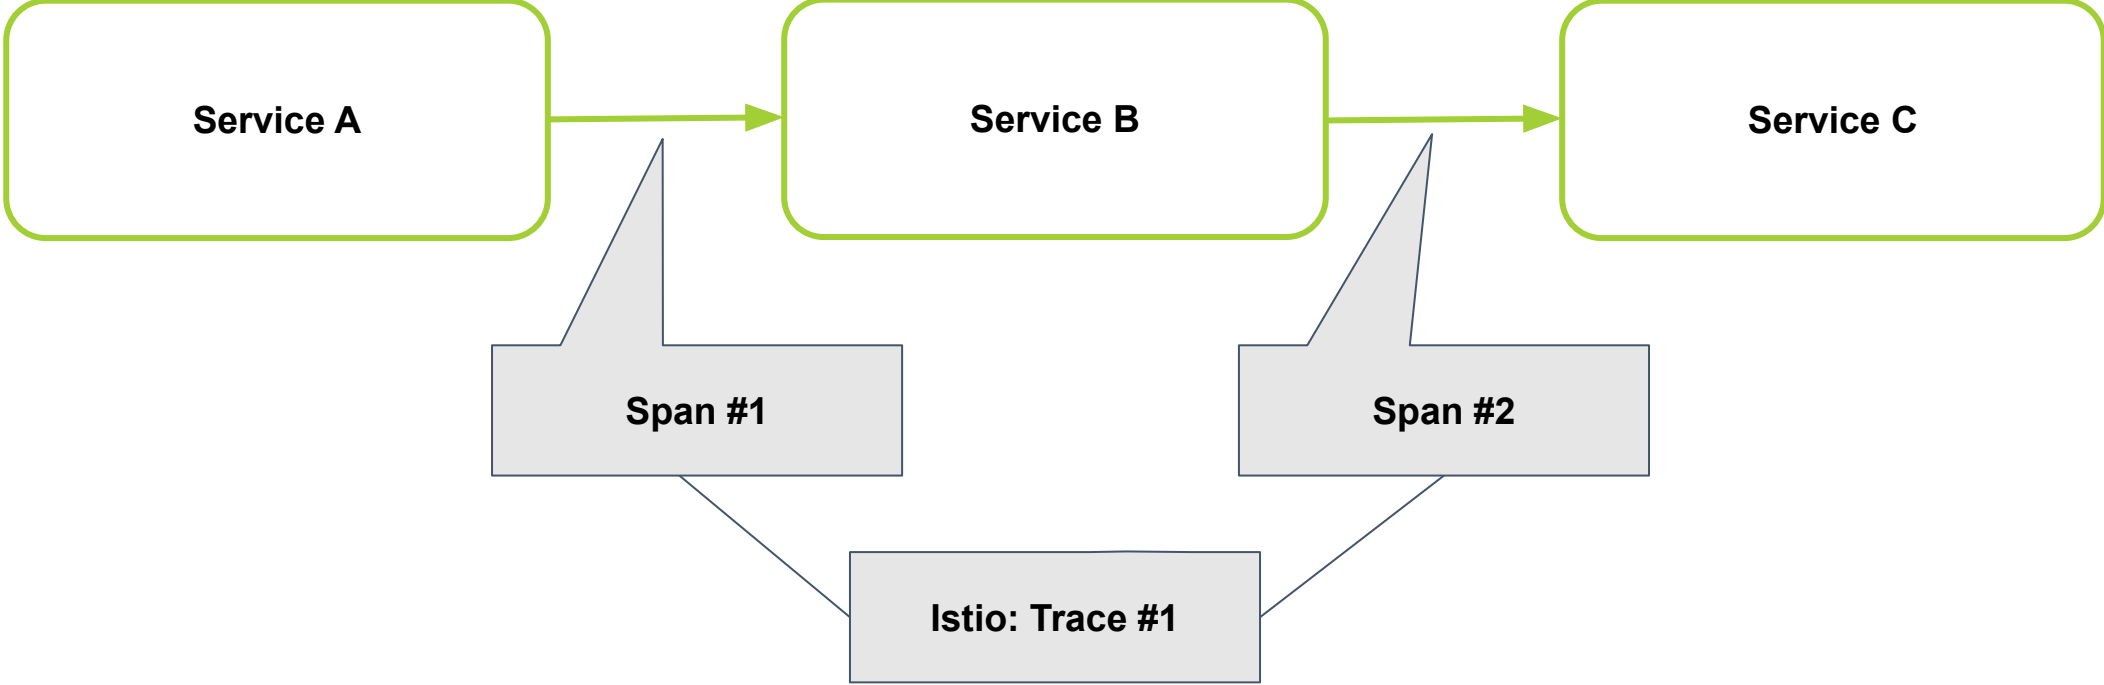
What we see of our own app



Without context propagation



With context propagation



The good news: Istio now supports OTEL

```
meshConfig:  
  extensionProviders:  
  - name: otel-tracing  
    opentelemetry:  
      port: 4317  
      service: otel-collector.svc.cluster.local
```

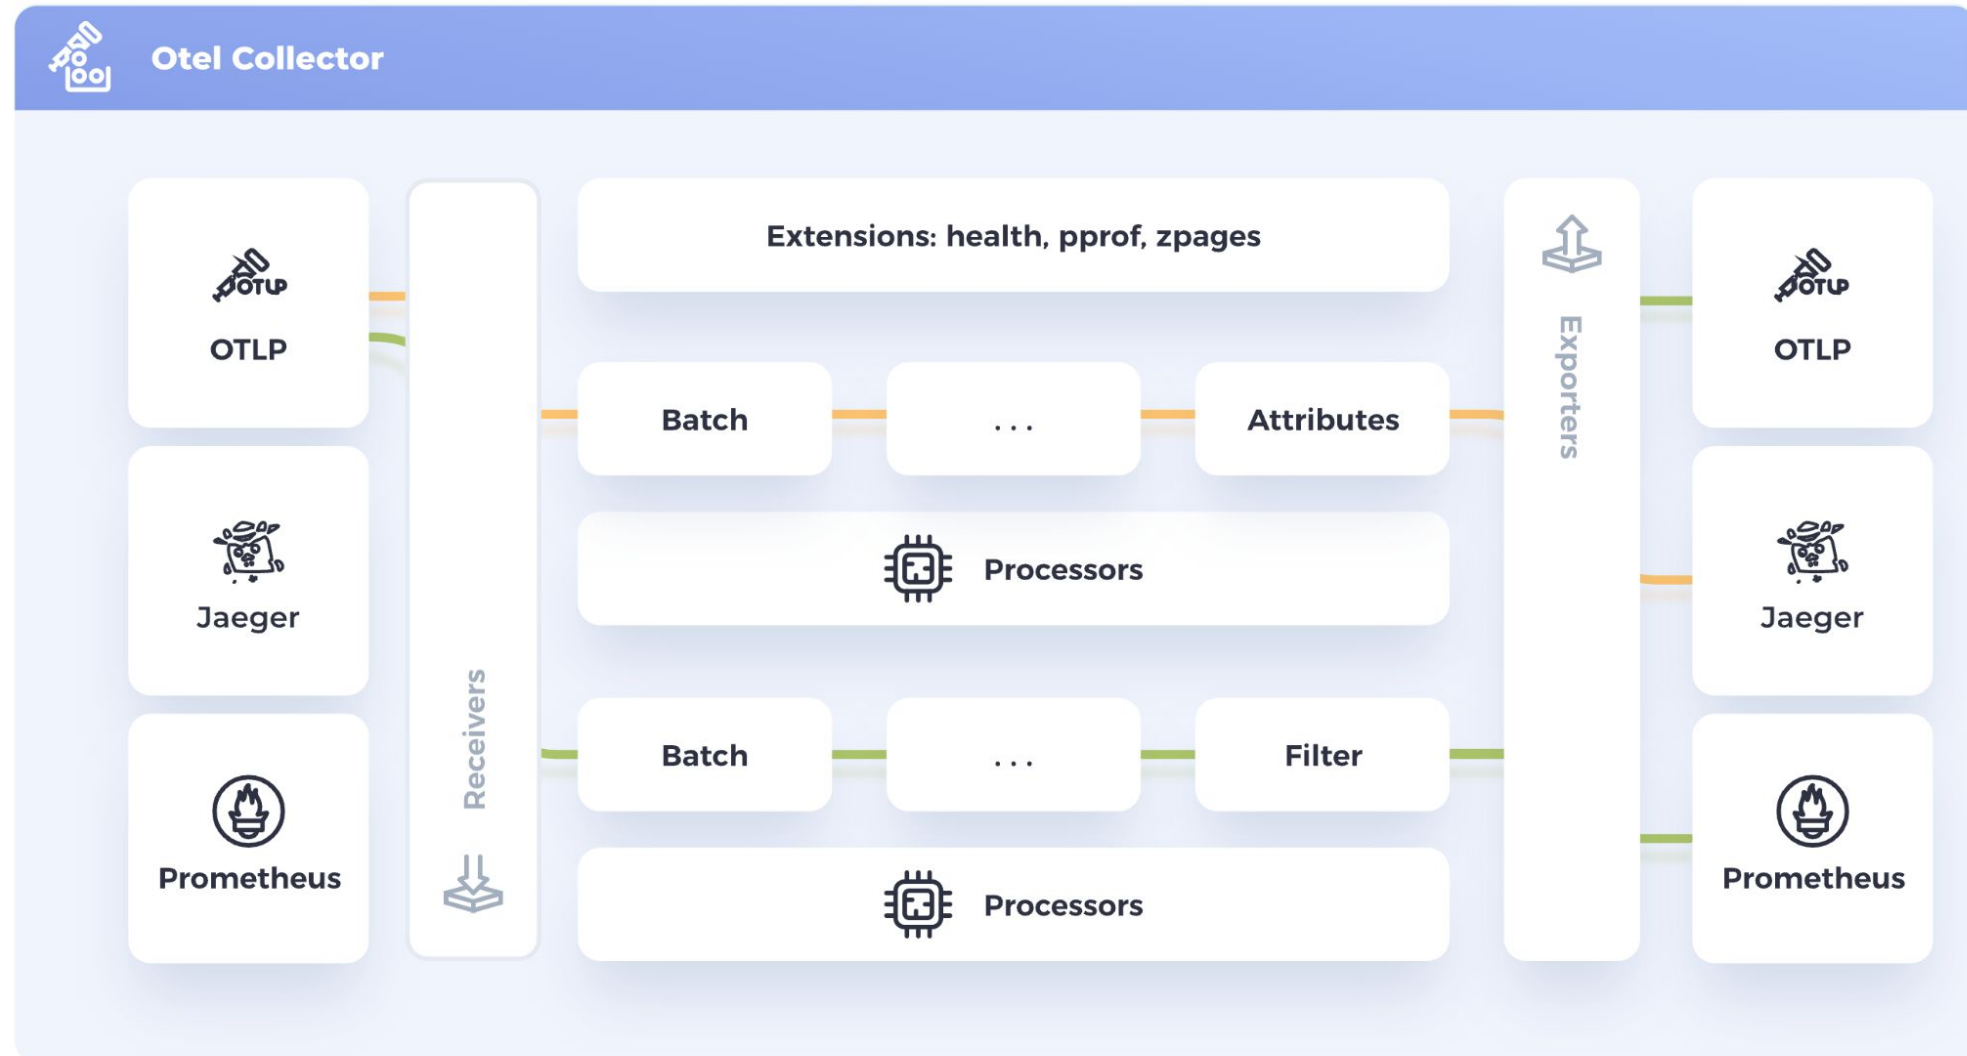
OpenTElemetry

- Framework and toolkit to create and manage telemetry data
- Can be used with a broad variety of observability backends

Components:

- **Collector** to receive, process and export telemetry data
- Code **Instrumentation** support for many languages

Otel Collector



Otel auto instrumentation

```
apiVersion: opentelemetry.io/v1alpha1
kind: Instrumentation
metadata:
  name: my-app-instrumentation
spec:
  propagators:
    - tracecontext
    - baggage
  sampler:
    type: always_on
  java:
    env:
      - name: OTEL_EXPORTER_OTLP_ENDPOINT
        value: http://otel-collector.otel.svc.cluster.local:4317
```

Considerations when getting started

- Istio is difficult, so is OpenTelemetry
- What backend to use?
- Use auto instrumentation or the SDK?
- Use Otel for everything (metrics, logs, traces)?
- Scaling of the Collector
- Sampling rates
- How to correlate telemetry data?
- How to offer a seamless experience to developers?

Developer and Operations User Experience

Self-service Portal
(otomi-console)

CLI
(otomi-shell)

Desired State Store
(otomi-values)

Golden Templates Catalog
(otomi-charts)

Control Plane (otomi-api)

Core (otomi-core)

Source Control

Pipelines

Artifacts

Continuous Deployment

Security

Proxy & Service Mesh

Observability

Storage & Backup

FinOps



Code

Ship

Run

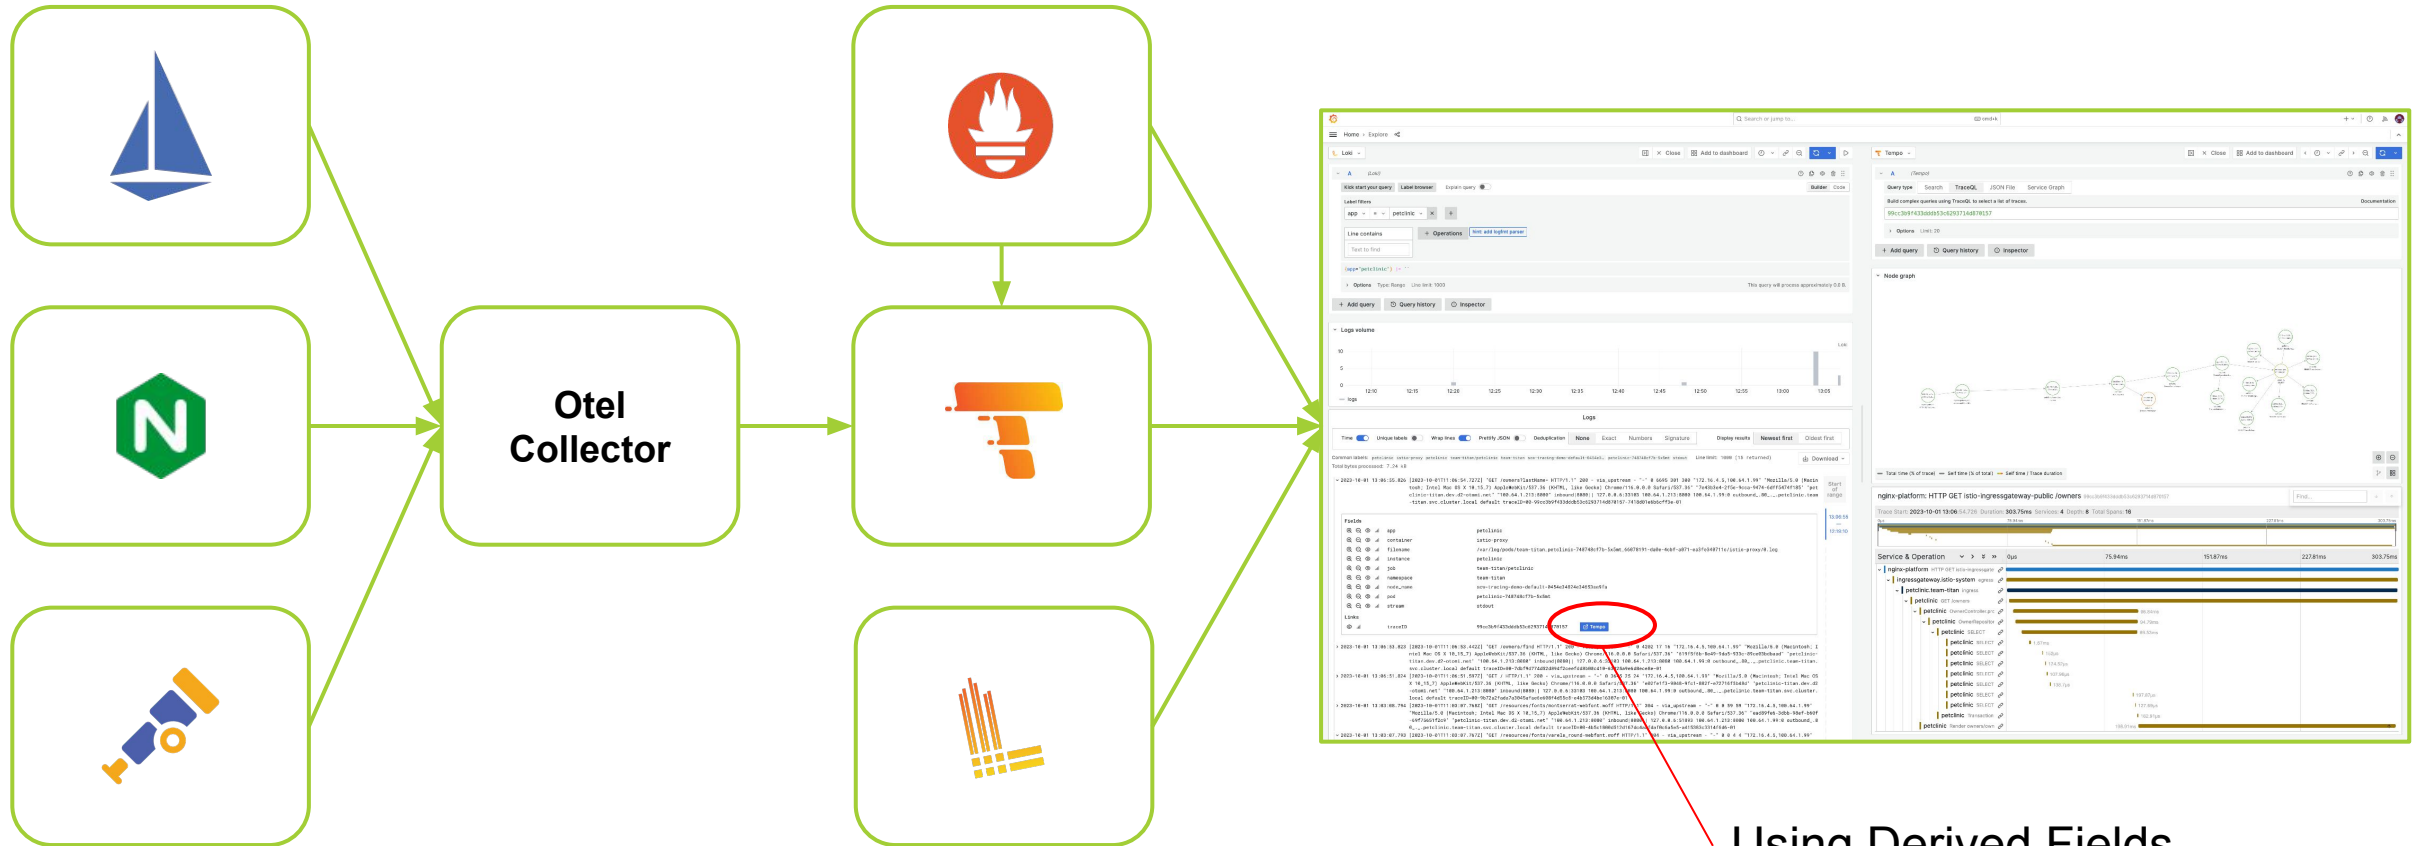
Automation (otomi-tasks)



Kubernetes



Observability in Otomi



Using Derived Fields

Otomi on GitHub

<https://github.com/redkubes/otomi-core>

And if you like the otomi project, give it a





ENTERPRISE
DEVOPS
T E C H C O N