K8s Cluster Logging with ClickHouse® & OpenTelemetry

A forward-looking solution for complete Kubernetes Cluster Logging

Josh Lee
Open Source Developer Advocate
josh@altinity.com

Why Logs Matter

A humble log…

```
2024-07-01 09:35:34 231ms GET /home 200
```

Controller Class: All | Namespace: live | Ingress: All | Ingress Pod: All | Config Reloads

Overview

Requests (period)	% Succe...	Conns (...	Reqs (2...	% Success (2m)	HTTP 1/2xx (2m)	HTTP 3xx (2m)	HTTP 4xx (2m)	HTTP 5xx (2m)
67.3 K	97.0%	311	850	99.3%	592	228	30	0

HTTP Requests / Ingress

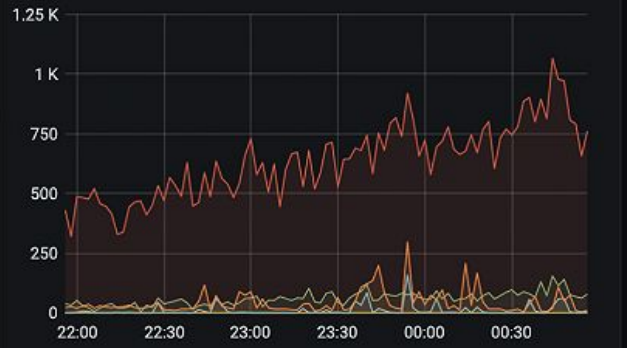

HTTP Status Codes

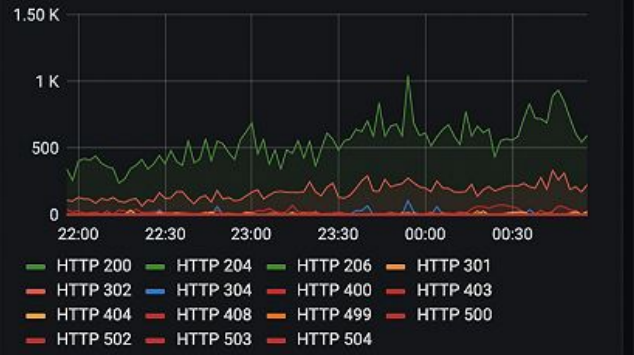

Total HTTP Requests

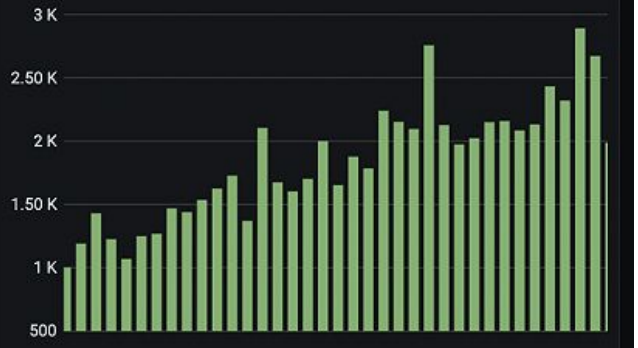

Latency

Latency (Average Percentiles)

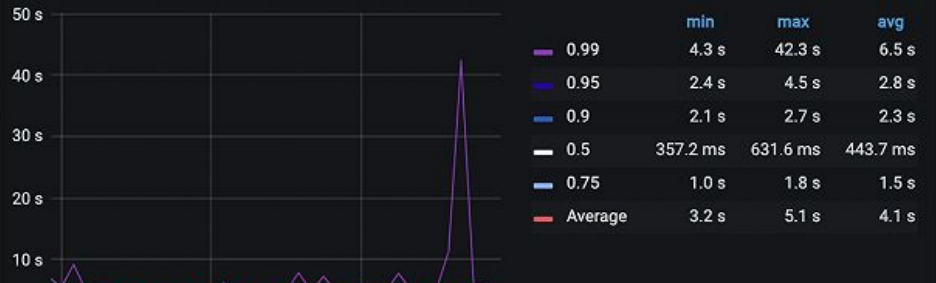

	min	max	avg
0.99	4.3 s	42.3 s	6.5 s
0.95	2.4 s	4.5 s	2.8 s
0.9	2.1 s	2.7 s	2.3 s
0.5	357.2 ms	631.6 ms	443.7 ms
0.75	1.0 s	1.8 s	1.5 s
Average	3.2 s	5.1 s	4.1 s

Latency Heatmap

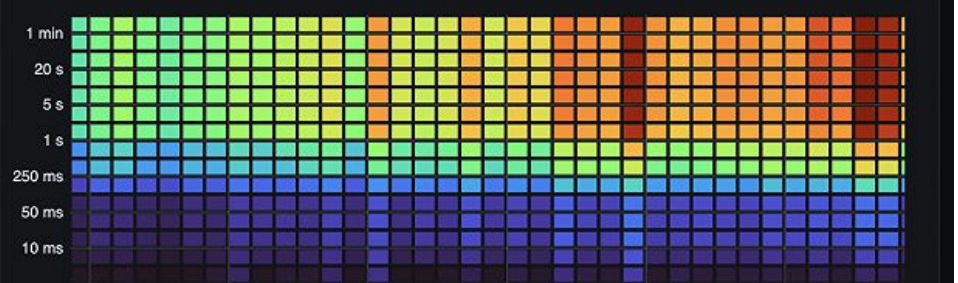

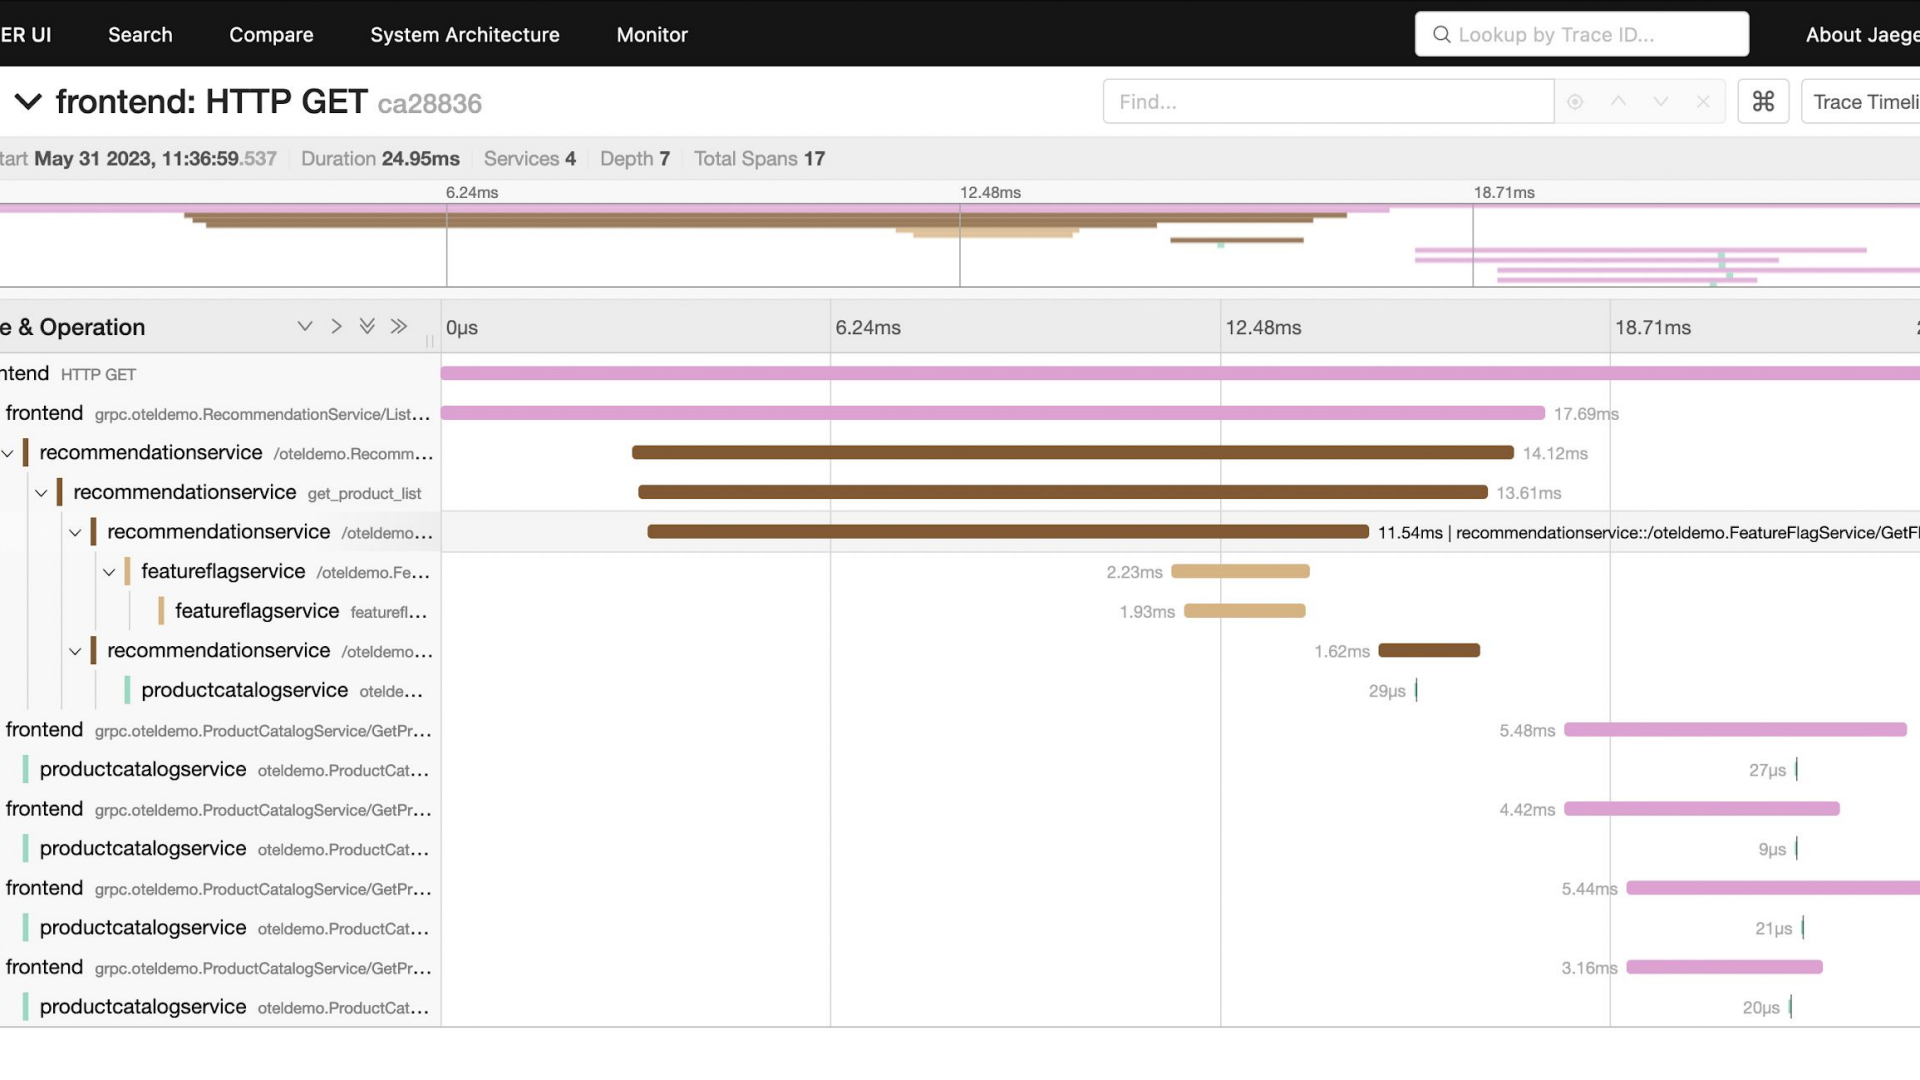
ER UI
Search
Compare
System Architecture
Monitor
Lookup by Trace ID...
About Jaege
frontend: HTTP GET ca28836
Find...
Trace Timeli
tart May 31 2023, 11:36:59.537
Duration 24.95ms
Services 4
Depth 7
Total Spans 17
6.24ms
12.48ms
18.71ms
e & Operation
0µs
6.24ms
12.48ms
18.71ms
ntend HTTP GET
frontend grpc.oteldemo.RecommendationService/List...
17.69ms
recommendationservice /oteldemo.Recomm...
14.12ms
recommendationservice get_product_list
13.61ms
recommendationservice /oteldemo...
11.54ms | recommendationservice::/oteldemo.FeatureFlagService/GetF
featureflagservice /oteldemo.Fe...
2.23ms
featureflagservice featurefl...
1.93ms
recommendationservice /oteldemo...
1.62ms
productcatalogservice otelde...
29µs
frontend grpc.oteldemo.ProductCatalogService/GetPr...
5.48ms
productcatalogservice oteldemo.ProductCat...
27µs
frontend grpc.oteldemo.ProductCatalogService/GetPr...
4.42ms
productcatalogservice oteldemo.ProductCat...
9µs
frontend grpc.oteldemo.ProductCatalogService/GetPr...
5.44ms
productcatalogservice oteldemo.ProductCat...
21µs
frontend grpc.oteldemo.ProductCatalogService/GetPr...
3.16ms
productcatalogservice oteldemo.ProductCat...
20µs

Observability is not any one signal

Metrics

Aggregable

Is there a problem?

Traces

Request-Scoped

Where is the problem?

Logs

Verbose, time-stamped records

What is the problem?

Why do we need application logs?

- Granular insights into system operations
- Root cause analysis and troubleshooting
- Incremental adoption of advanced observability
- Rich data for AI models

Logging in Kubernetes

Logging in Kubernetes

1. No built-in cluster-wide logging
2. No standard log formats
3. Contextual disconnection
4. Containers are ephemeral
5. $$$

Cluster Logging is Essential

1. Retention and Querying
2. Continuity
3. Annotations, normalization, filtering
4. Combined view of applications & infrastructure
5. Centralized scalability
6. Advanced analytics & data-mining

The OpenTelemetry Collector

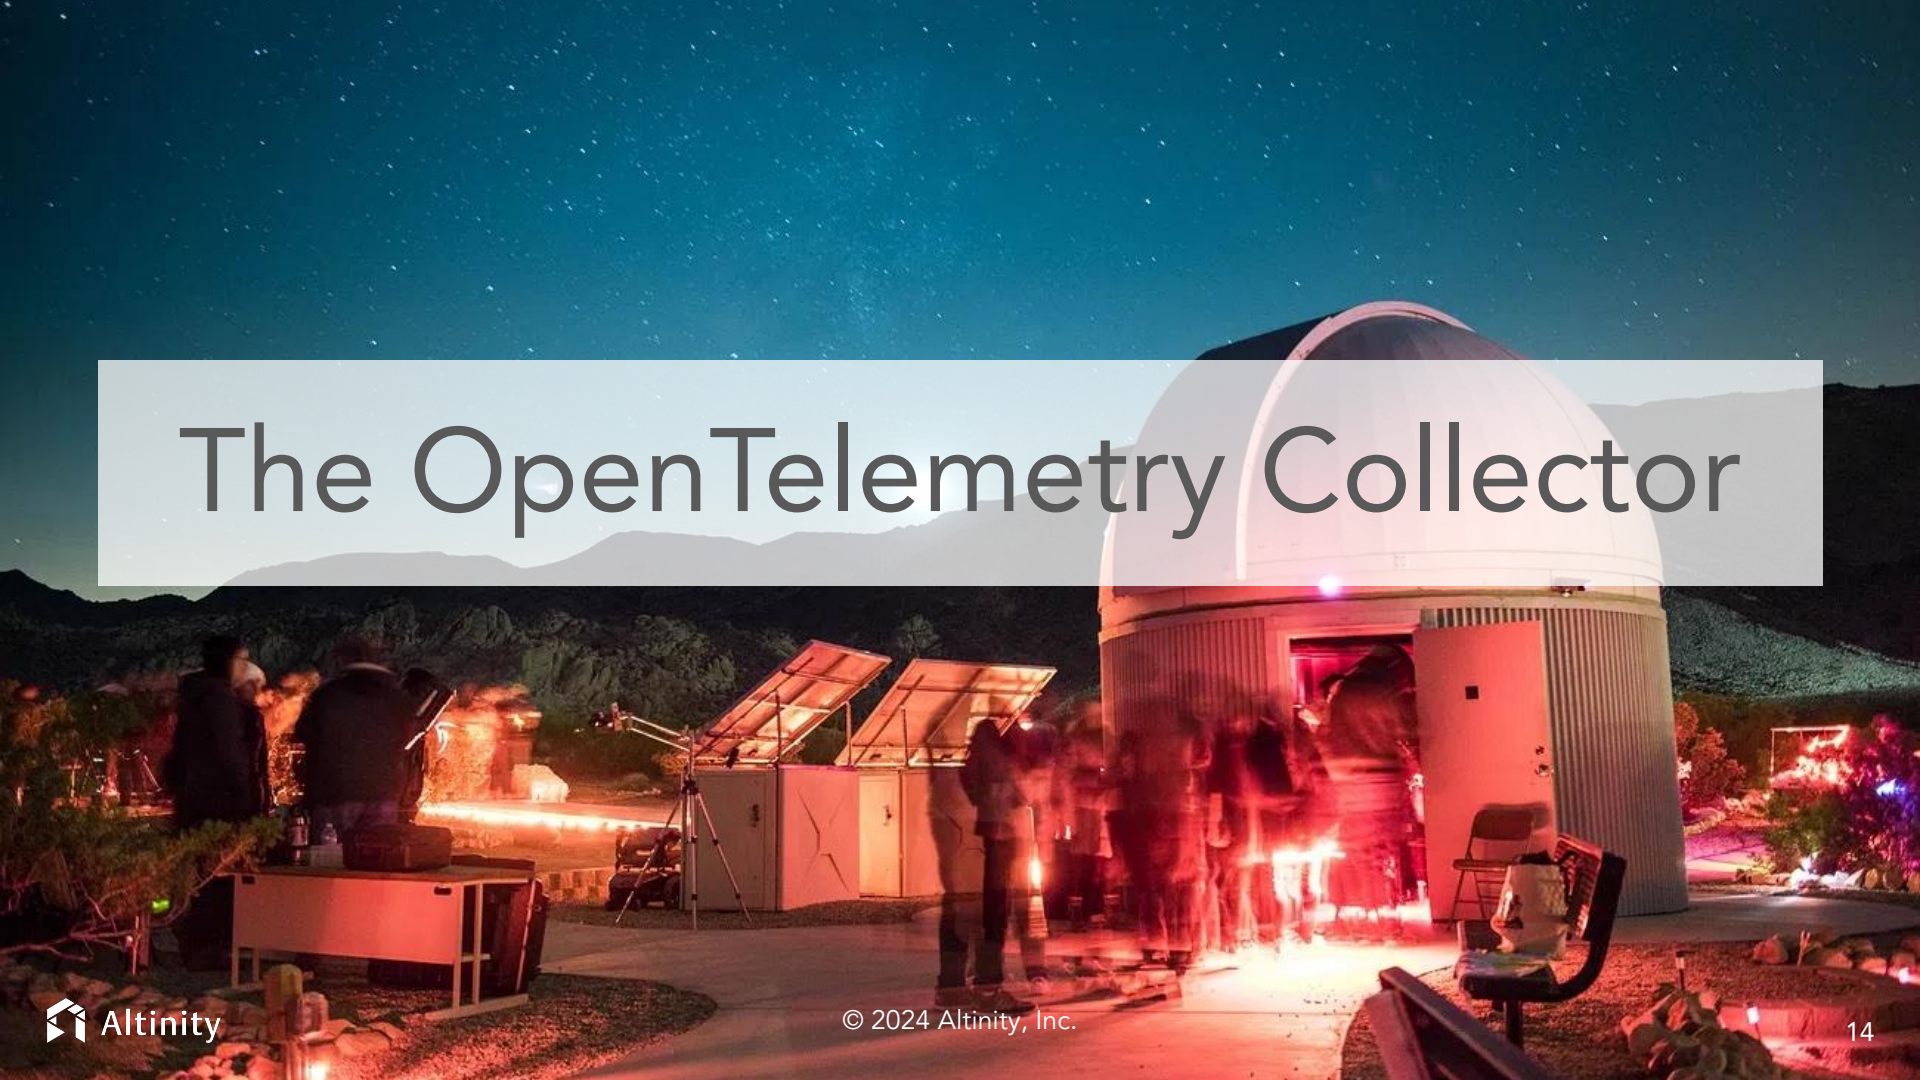

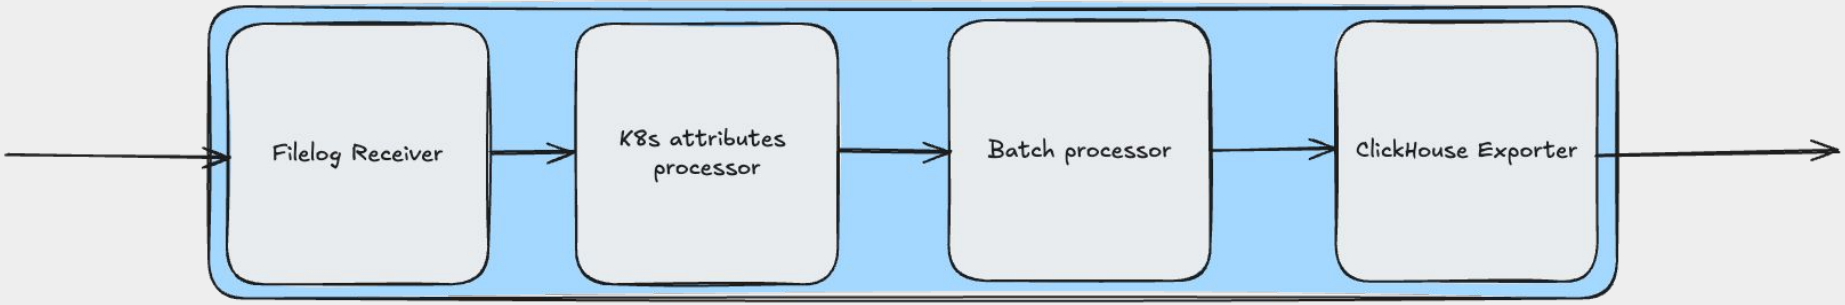

OpenTelemetry Collector Pipeline

Resource Attributes

Extensions for Kubernetes

- Kubernetes Attributes Processor
- Filelog Receiver
- Kubelet Metrics Receiver

The Altinity Operator for ClickHouse®

Our Demo Logging Architecture: Kubernetes + ClickHouse® + OpenTelemetry

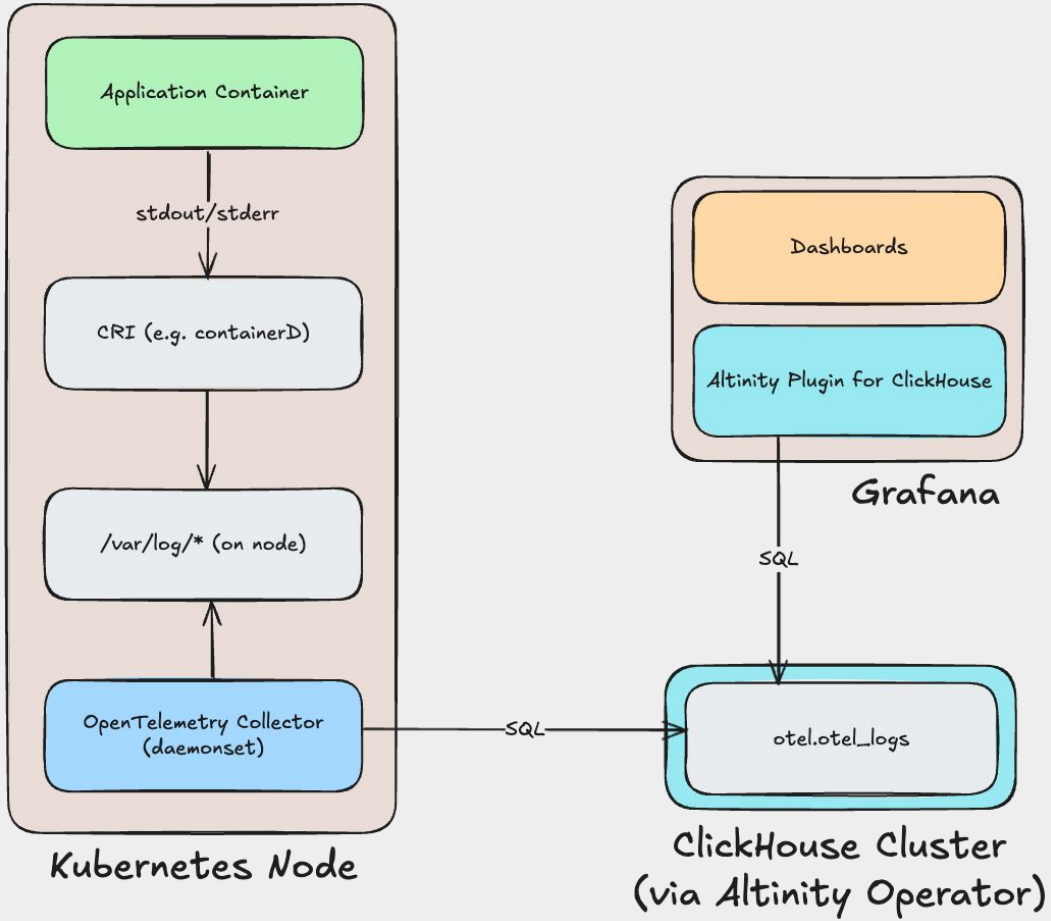

Why ClickHouse® for Log Storage?

Demo: Ingesting Logs with OpenTelemetry + ClickHouse

Demo Agenda

1. Setting Up the OpenTelemetry Collector
2. Configuring the Exporter to ClickHouse
3. Deploying ClickHouse® with the Altinity® Operator
4. Generating Logs from a Sample Application
5. Verifying Log Ingestion in ClickHouse
6. Querying and Visualizing Logs in Grafana®
7. Scaling for production
8. Q&A

Demo

Scaling the Solution

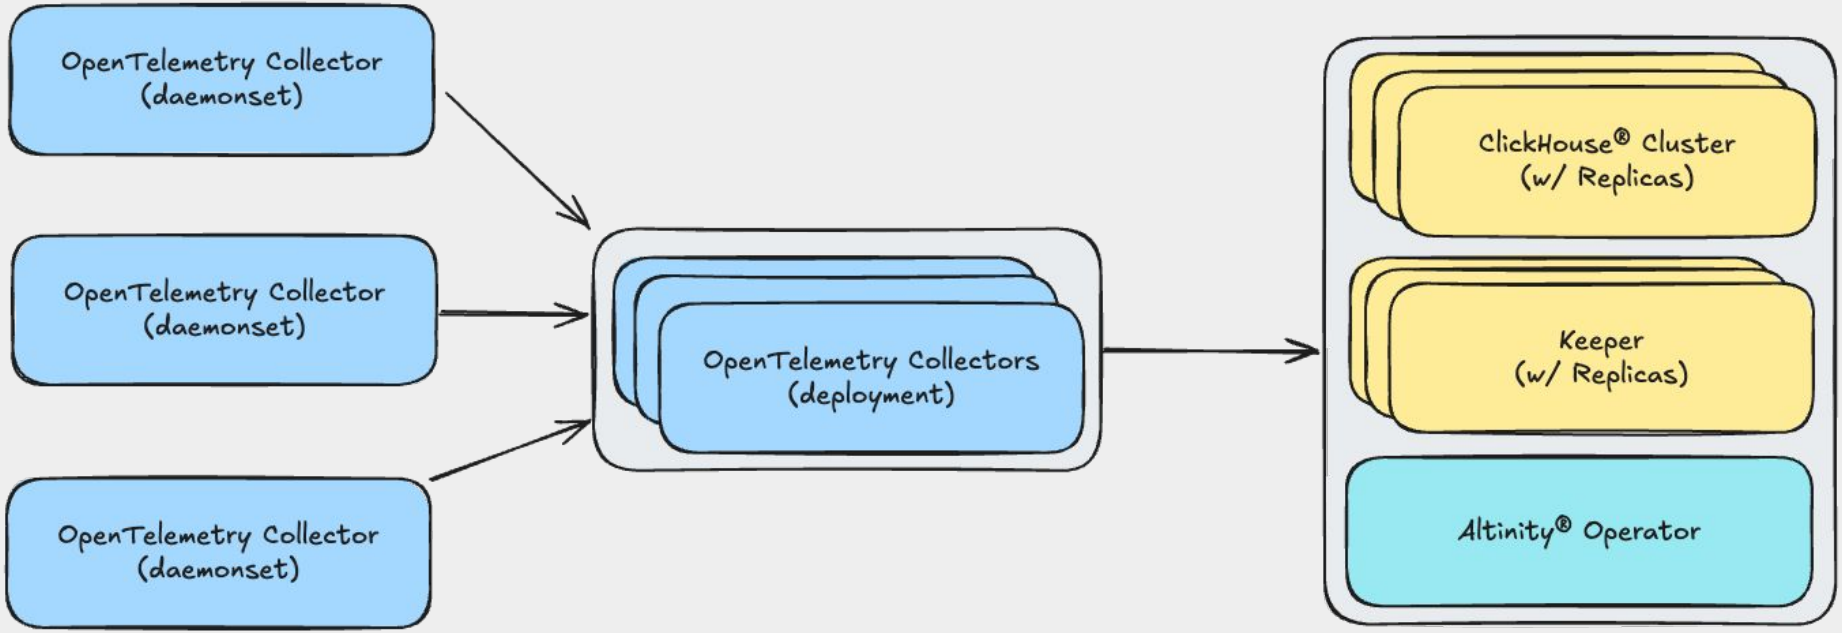

Q&A