

# Observability of Continuous Delivery Pipelines with OpenTelemetry: A Case Study with Spinnaker

Noel McKeown

School of Enterprise Computing and Digital Transformation, TU Dublin, Ireland

X00218705@mytudublin.ie

## Introduction

Continuous Delivery (CD) pipelines are the backbone of modern software release processes, yet many organisations lack meaningful visibility into how these pipelines perform. This absence of pipeline level observability creates a blind spot that limits operational awareness, impedes root cause analysis and restricts the ability to measure delivery performance objectively. This research investigates how OpenTelemetry (OTel), a vendor neutral observability technology can be extended beyond traditional application monitoring to instrument continuous delivery pipeline lifecycle events. Using the Spinnaker CD platform as a case study, the research demonstrates how native pipeline execution events can be captured, transformed into structured telemetry and exported as Prometheus compatible metrics. These metrics are further aligned with DORA performance indicators, enabling empirical measurement of deployment speed, reliability and recovery.

## Observability Gap in CD Pipelines

Continuous delivery platforms remain under instrumented, leaving pipeline health, performance and failure behaviour poorly measured. This research identifies a critical observability gap where pipeline health, execution duration, failure patterns and recovery behaviour are not consistently measurable. Without pipeline metrics, teams cannot proactively identify bottlenecks, assess release stability or support service level objectives. Addressing this gap is essential for data driven DevOps decision making.

## OTel Integration Architecture

A prototype architecture was implemented to integrate Spinnaker with OpenTelemetry by capturing pipeline lifecycle events via a lightweight webhook receiver and processing them through an OpenTelemetry Collector. Events are normalised into consistent metrics for ingestion by Prometheus and visualisation in Grafana using only open source components. The results demonstrate OpenTelemetry's viability as a unified observability layer for continuous delivery platforms.

## Mapping Pipeline Telemetry to DORA Metrics

A core contribution of this research is the mapping of raw pipeline execution events to DORA metrics. Deployment Frequency, Lead Time for Change, Change Failure Rate, and Mean Time to Restore. By extracting timestamps, execution states, trigger metadata and failure signals from pipeline telemetry, these metrics can be computed automatically in near real-time. Grafana dashboards and alerts provide actionable insights into delivery performance, enabling teams to correlate operational behaviour with business outcomes. This approach operationalises the pipeline metrics and provides a DORA scorecard for CD performance.

### Deployment Frequency & Lead Time for Change

Pipeline execution events provide precise timestamps for when a deployment starts and successfully completes. These signals enable Deployment Frequency to be calculated directly from successful pipeline runs over time, while Lead Time for Change is derived from execution start to completion duration.

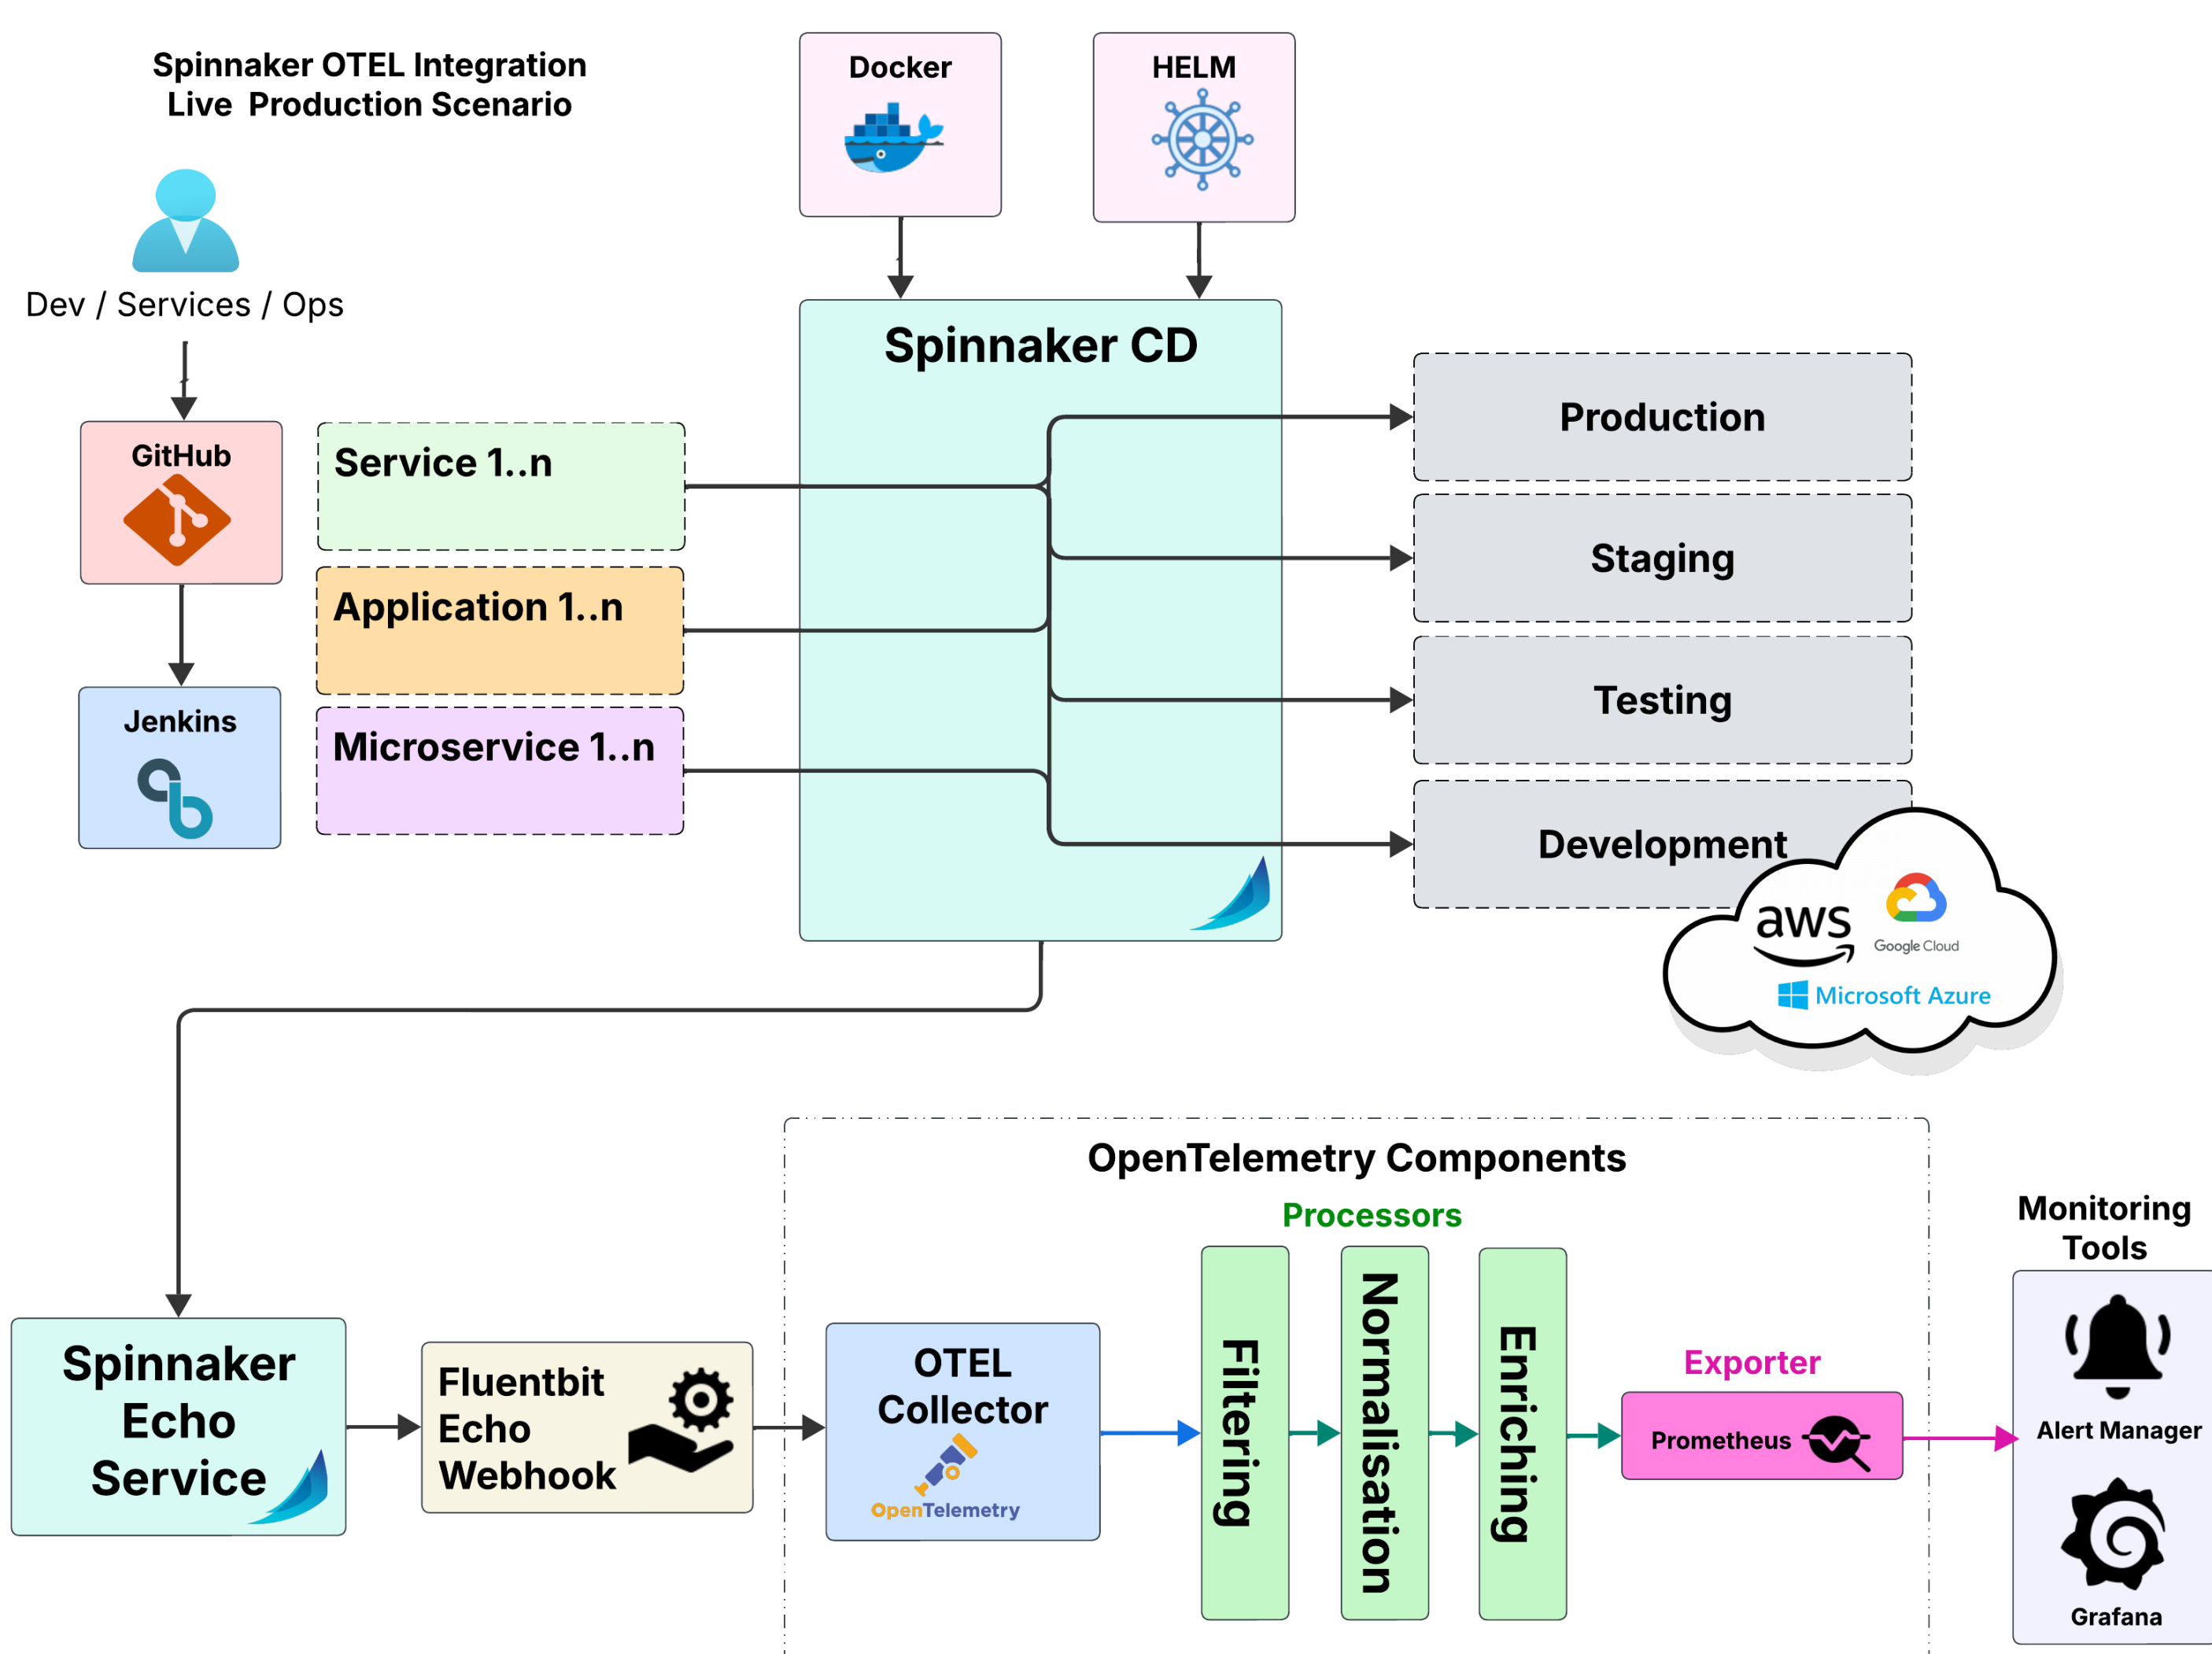
### Change Failure Rate & Reliability Signals

Pipeline failure and cancellation events are transformed into structured metrics that quantify Change Failure Rate. Failed executions, rollbacks and aborted stages provide a direct signal of deployment instability and release risk.

### Mean Time to Restore & Operational Recovery

Recovery related pipeline events enable calculation of Mean Time to Restore (MTTR) by measuring the interval between a failure event and the next successful execution. This captures how quickly teams can recover service using pipeline driven remediation or rollback workflows.

## Architecture and DORA Dashboard



## Conclusions and Future Work

Extending OpenTelemetry into CD pipelines is feasible and valuable, providing unified observability, actionable pipeline insights and objective DevOps performance measurement through an open source repeatable architecture.

Future work should validate the approach at enterprise scale, optimise telemetry performance, apply security standards and extend pipeline observability to other CD platforms and OpenTelemetry conventions.

## QR Code for Recording

