



Faster, Cheaper, Smarter: Lessons from Building Zero-Instrumentation Observability with Fluent Bit OpenSearch & Prometheus



June 5th, 2025

Hello!



Nilushan Costa

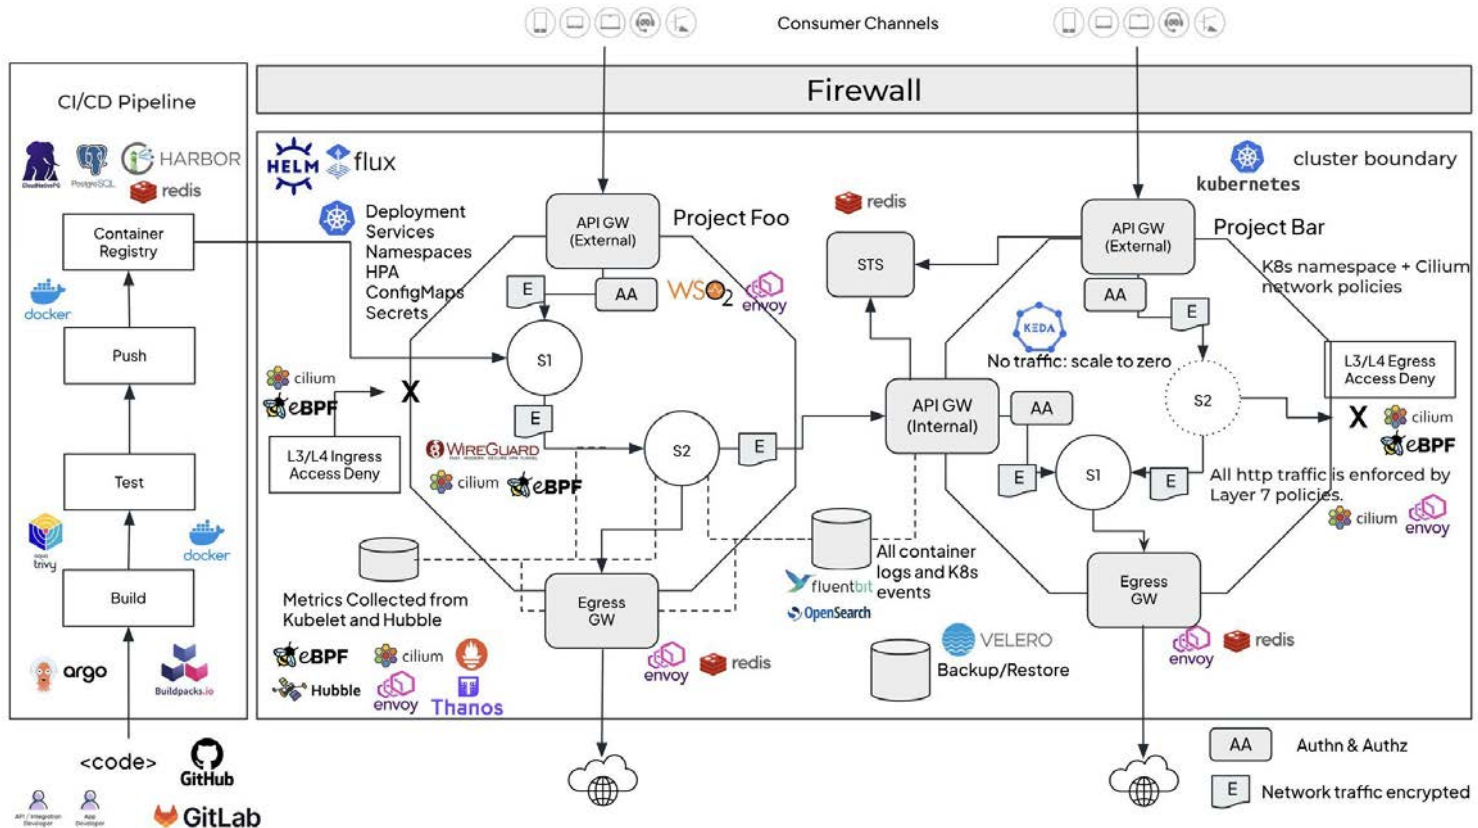
Associate Technical Lead



Isala Piyarisi

Senior Software Engineer

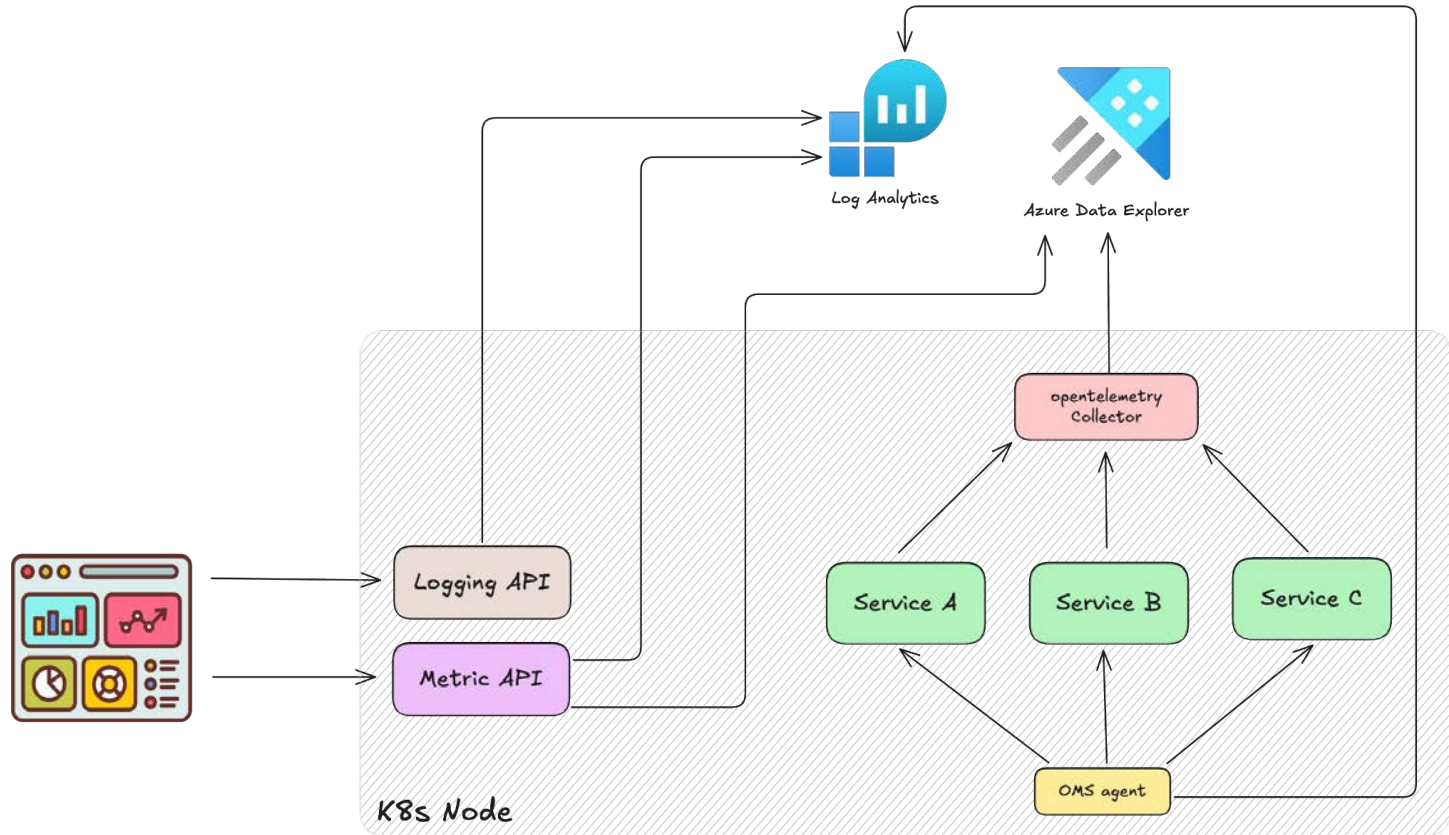
Who We Are & What is Choreo?



History




The "Before" State: Azure-centric Observability



The Mounting Challenges



Skyrocketing Costs



Multi cloud



Limited Control & Flexibility

Our Goals: A User-Centric, Future-Proof Vision

- For Choreo Users:

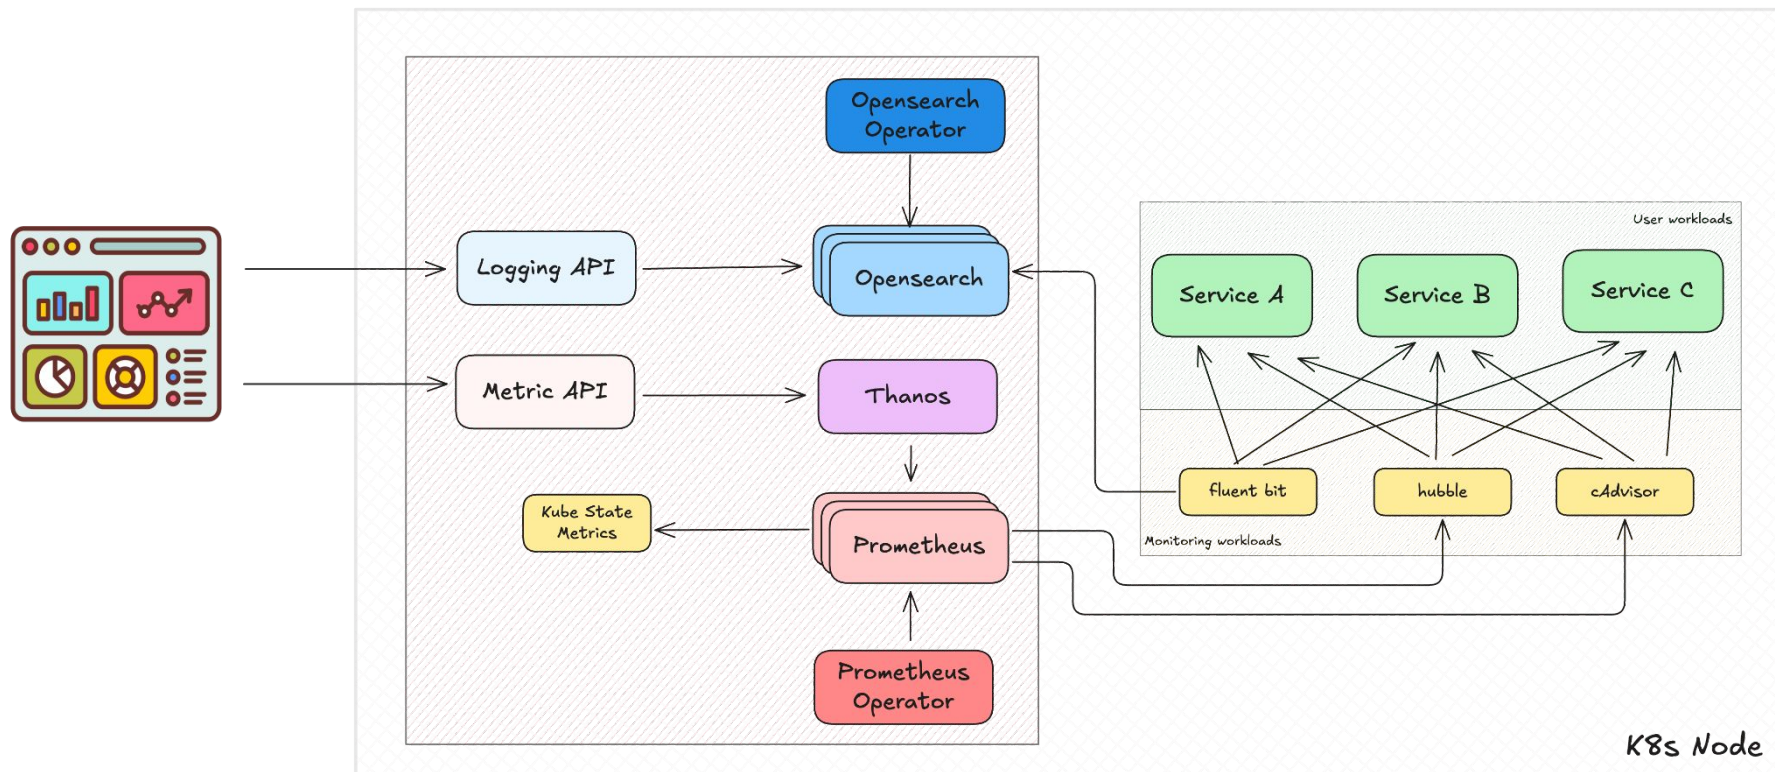
- Speed
- Smarter Tools
- Zero-Instrumentation

- For the Choreo Platform:

- Cost Reduction
- Cloud Agnostic
- Data Sovereignty
- Control & Innovation



The New Stack: OSS Power!



Logs



Solving the Log Latency Problem

- *“The average latency to ingest log data is between 20 seconds and 3 minutes. The specific latency for any particular data will vary depending on several factors”*
- Goal: ≤ 5 seconds
- “tail -f” experience
- Fluent Bit + OpenSearch latency



Solving the Log Latency Problem

- OpenSearch as a cache
- Combine data from 2 APIs
 - ⦿ Live logs API - OpenSearch
 - ⦿ Historical logs API - Log Analytics



Logs Stack

- Fluent Bit
- OpenSearch
- OpenSearch Dashboard
- Logging API



Deployment

- Helm chart based on upstream chart
- Distributed deployment of OpenSearch
- OpenSearch Operator
- Multiple layers of backups



Cost Savings

- ~USD 1800 per month saving for one dataplane in Azure
- ~USD 2500 per day for one customer in AWS

Lessons

- Store raw logs if schema is not the same
- Log throttling
- Health monitoring
- Storage medium
- Look at the stack as a whole when debugging
- Fluent Bit Inotify_Watcher



Logs

Runtime Logs

Filter

Time: Past 10 minutes

Type

Environment

Deployment Track

Time Range

> 2025-05-28T10:34:31.145+05:30 Development Application Logs v1.0 unhandled error returned from the service module="ballerina/http" error={"causes":[],"message":"Error","detail":{},"stackTrace":[{"callableName":"\$get\$fa

> 2025-05-28T10:34:32.127+05:30 Development Gateway Logs v1.0 500 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/failure" ServicePath="/srvc/failure" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

> 2025-05-28T10:34:35.991+05:30 Development Application Logs v1.0 {"a":"valueofa","b":"valueofb"} module="nilushan/success_failure"

> 2025-05-28T10:34:36.829+05:30 Development Application Logs v1.0 {"a":"valueofa","b":"valueofb"} module="nilushan/success_failure"

> 2025-05-28T10:34:36.892+05:30 Development Gateway Logs v1.0 202 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/jsonlog" ServicePath="/srvc/jsonlog" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

> 2025-05-28T10:34:38.138+05:30 Development Gateway Logs v1.0 202 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/jsonlog" ServicePath="/srvc/jsonlog" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

< 2025-05-28T10:35:25.868+05:30 Development Application Logs v1.0 Request received at /success endpoint module="nilushan/success_failure"

timestamp

2025-05-28T05:05:25.868Z

level

INFO

logLine

Request received at /success endpoint

logContext

{

"module" : "nilushan/success_failure"

}

componentVersion

v1.0

componentVersionId

d52a44e1-7a38-43f7-810c-fe6ae7be2d56

envName

Development

> 2025-05-28T10:35:26.784+05:30 Development Application Logs v1.0 Request received at /success endpoint module="nilushan/success_failure"

> 2025-05-28T10:35:26.992+05:30 Development Gateway Logs v1.0 200 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/success" ServicePath="/srvc/success" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

> 2025-05-28T10:35:27.887+05:30 Development Application Logs v1.0 Request received at /success endpoint module="nilushan/success_failure"

> 2025-05-28T10:35:28.238+05:30 Development Gateway Logs v1.0 200 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/success" ServicePath="/srvc/success" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

> 2025-05-28T10:35:28.996+05:30 Development Gateway Logs v1.0 200 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/success" ServicePath="/srvc/success" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

> 2025-05-28T10:35:45.414+05:30 Development Application Logs v1.0 Request received at /success endpoint module="nilushan/success_failure"

> 2025-05-28T10:35:46.275+05:30 Development Gateway Logs v1.0 200 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/success" ServicePath="/srvc/success" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

> 2025-05-28T10:35:46.296+05:30 Development Application Logs v1.0 Request received at /success endpoint module="nilushan/success_failure"

> 2025-05-28T10:35:48.278+05:30 Development Gateway Logs v1.0 200 SERVICE_RESPONSE Method="GET" RequestPath="/q9cf/success-failure/v1.0/success" ServicePath="/srvc/success" UserAgent="Mozilla/5.0 (Macintosh; Intel Mac OS

> 2025-05-28T10:35:50.071+05:30 Development Application Logs v1.0 Request received at /failure endpoint module="nilushan/success_failure"

> 2025-05-28T10:35:50.073+05:30 Development Application Logs v1.0 unhandled error returned from the service module="ballerina/http" error={"causes":[],"message":"Error","detail":{},"stackTrace":[{"callableName":"\$get\$fa


Metrics




High level goals



HTTP Metrics



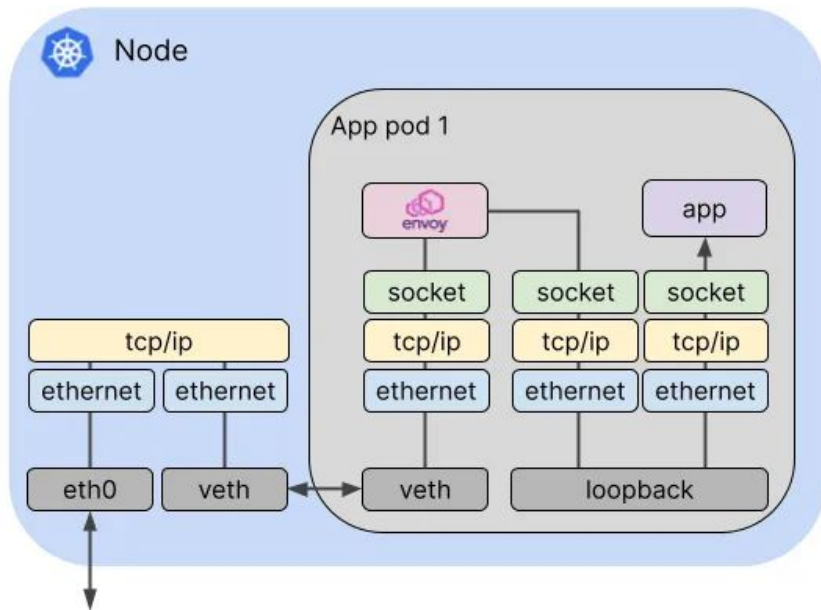
Usage Metrics



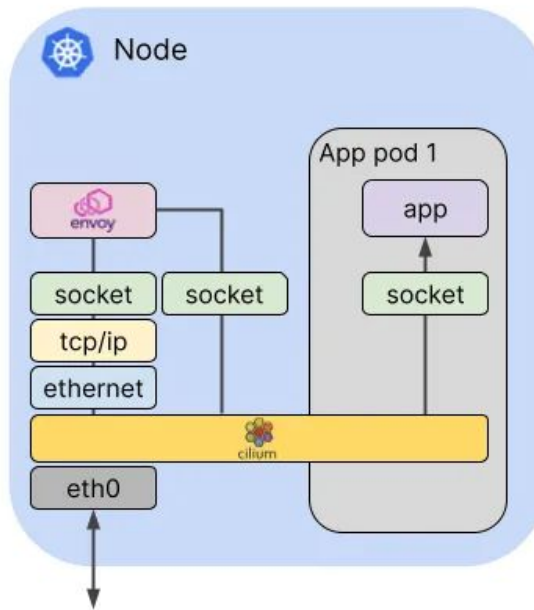
Service diagram

Deep Dive: HTTP Metrics via Hubble

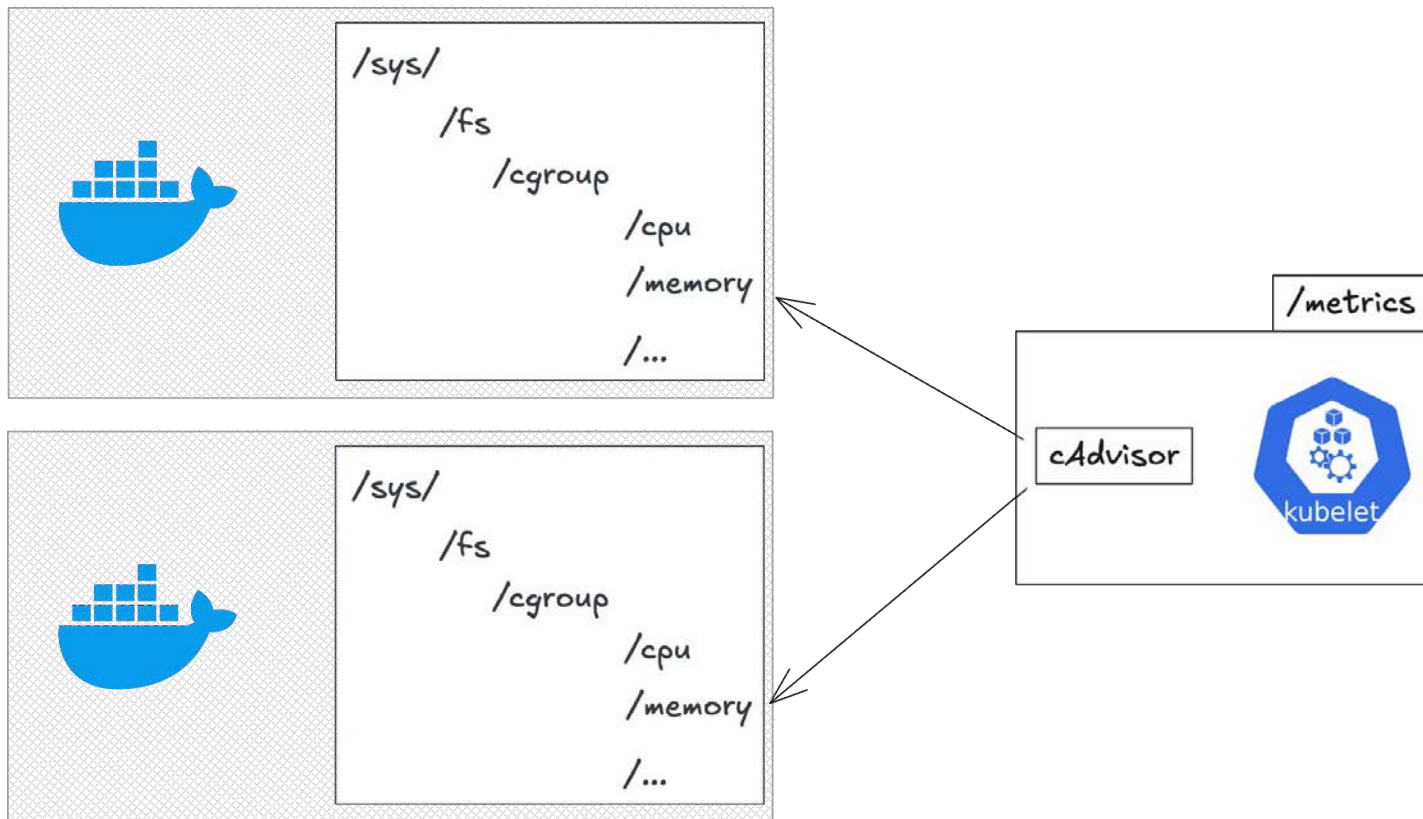
Service mesh with traditional networking



Sidecarless model, eBPF acceleration



Deep Dive: Usage via cAdvisor




Deep Dive: Kubernetes Context



Deep Dive: Reliability



Thanos



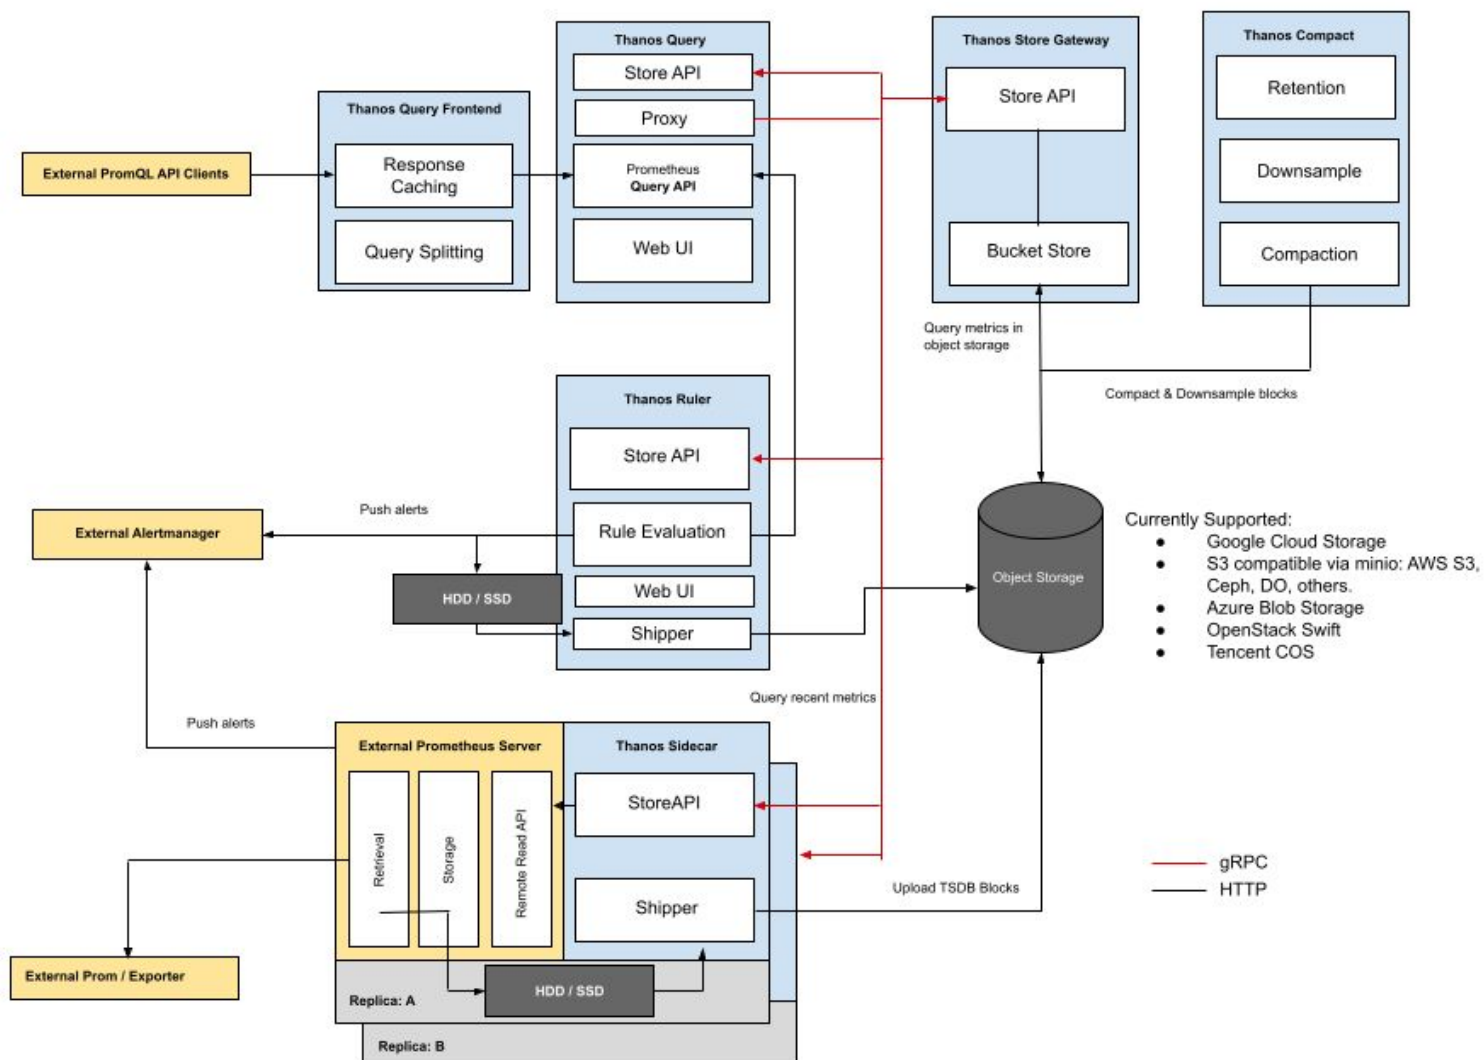
AWS S3



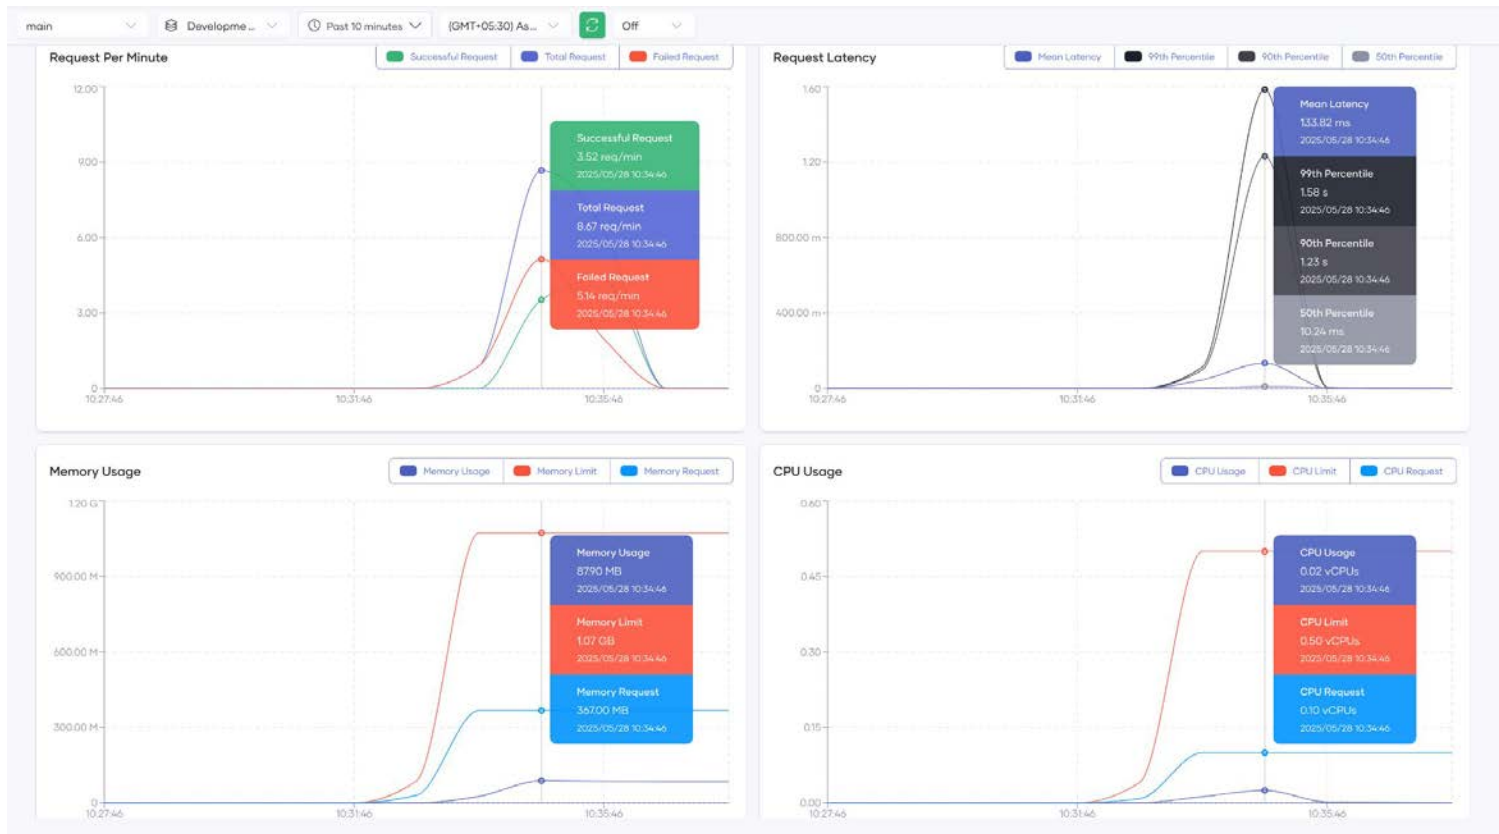
Azure Blob Storage



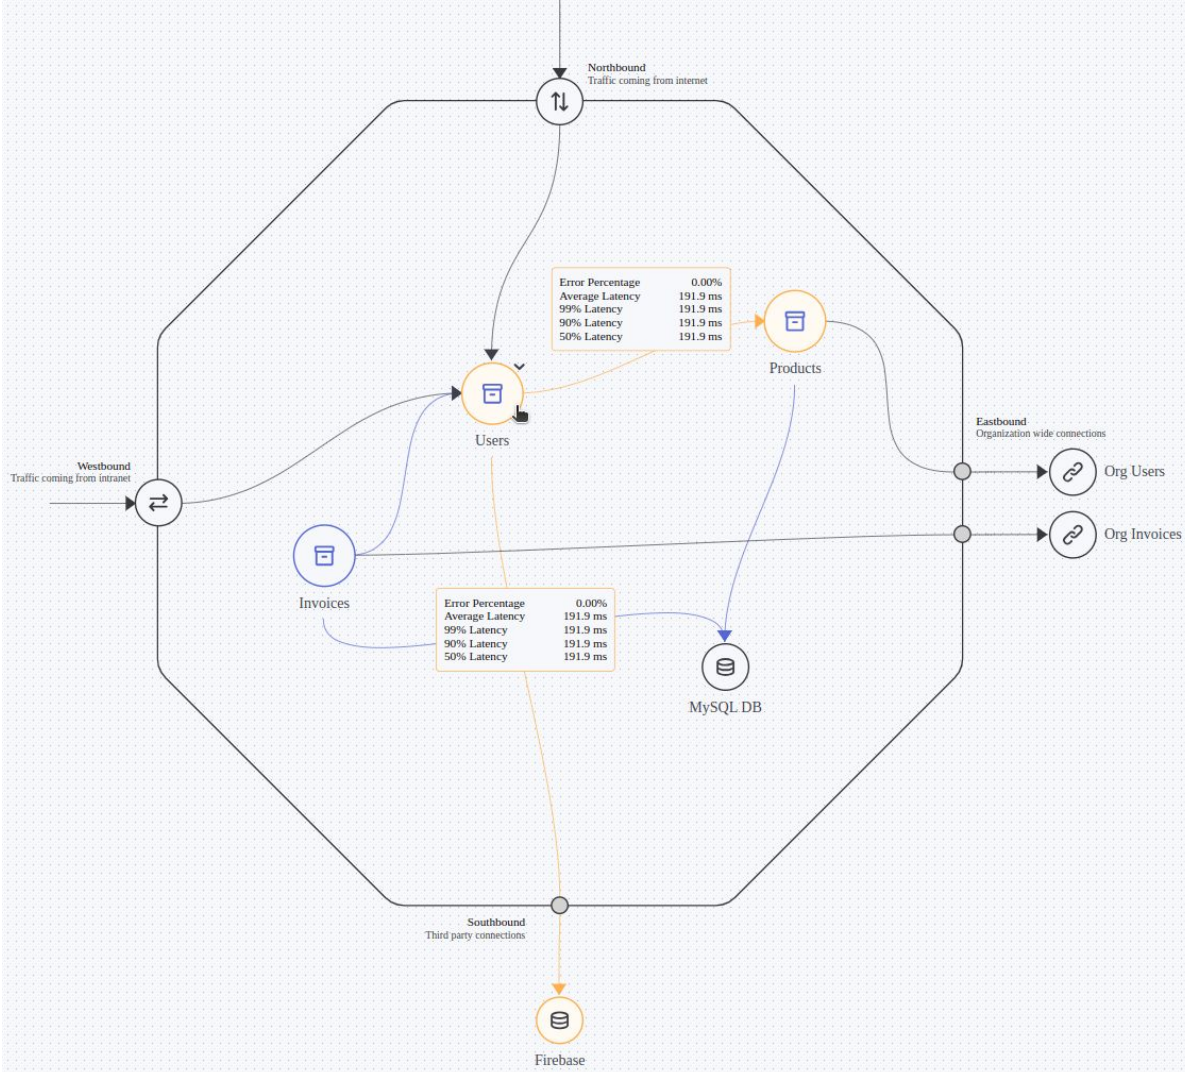
GCP Cloud Storage



Metrics



Metrics



Learned the hard way

- If you want accuracy use logs
- If there is an operator use the operator
- /metric cleanups are very important
- NFS are a bad idea
- $\text{sum of rates} \neq \text{rate of sums}$



Cost wise

- 1200 USD for ADX
- 1000 USD for Log Analytics
- ~250 USD in Compute
- ~100 USD in Storage

Conclusion



Conclusion

- Ability to easily enable Observability for new providers
- Ease of adding new features on top of the existing stack
 - ⦿ Alerting
- Ease of debugging – Same setup everywhere
- Cost reduction



Thanks!



wso2.com

