


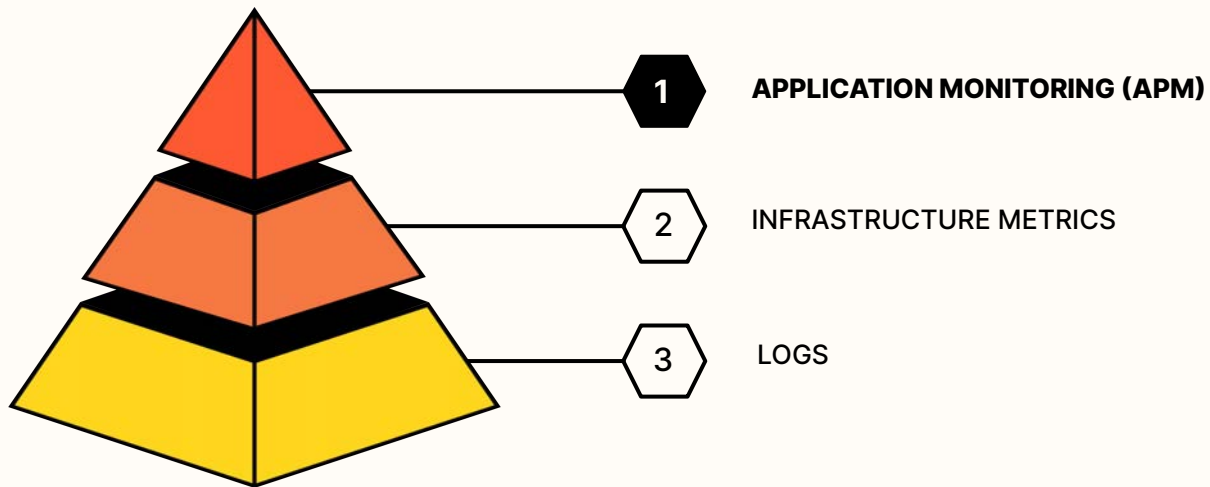
groundcover

# Cloud-Native Observability

True Kubernetes Observability with eBPF

Conf42 Cloud Native March 2023





**Observability is a Core Competency of any Team**



## O11y Spend are **20-30%** of Infra Spend (!)

Charity Majors, Honeycomb.io

## Ingest

STARTING AT

\$ **0.10**

Per ingested or scanned GB, per month\*

Ingest, process, live tail, and archive all logs

- Enrich and structure log data
- Parse on ingestion
- Generate log-based metrics
- Self-hosted archives, with the option to rehydrate
- Dynamic index routing

\*Per GB of uncompressed data ingested for processing, or compressed data scanned for rehydrating.

START FREE TRIAL

## Retain or Rehydrate

15-DAY RETENTION ▾

\$ **1.70**

Per million log events per month\*

Retain logs based on their value and rehydrate from archives on-demand


- Define log retention based on tags or facets
- Simplified pricing based on retention for better cost control
- Log patterns and analytics
- Log Rehydration™ for audits and historical analysis

\*Billed annually or \$2.55 on-demand

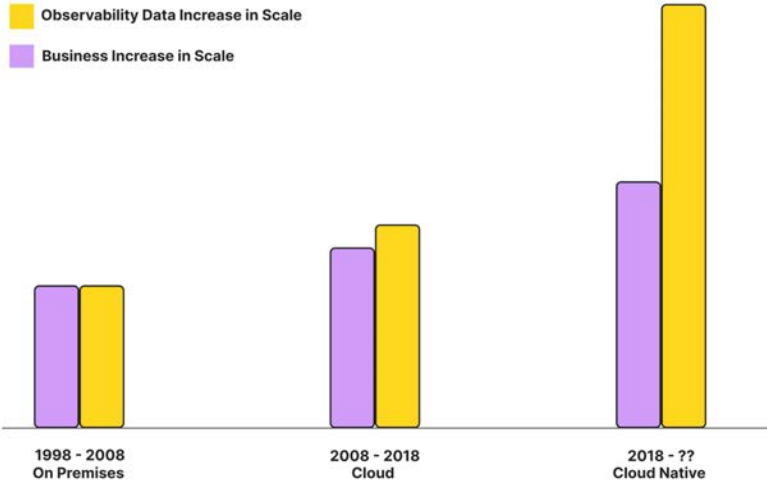
START FREE TRIAL

# But mostly, Unpredictable!

Datadog Pricing, 2023




## Observability data is growing 2-3x faster than application data



# Cloud-native Made Things Worse


Cloud Native Monitoring, O'REILLY




**“Most companies have 184  
microservices on average”**

Survey by Kong, 2022

Observability was built  
**to be part of a dev cycle**



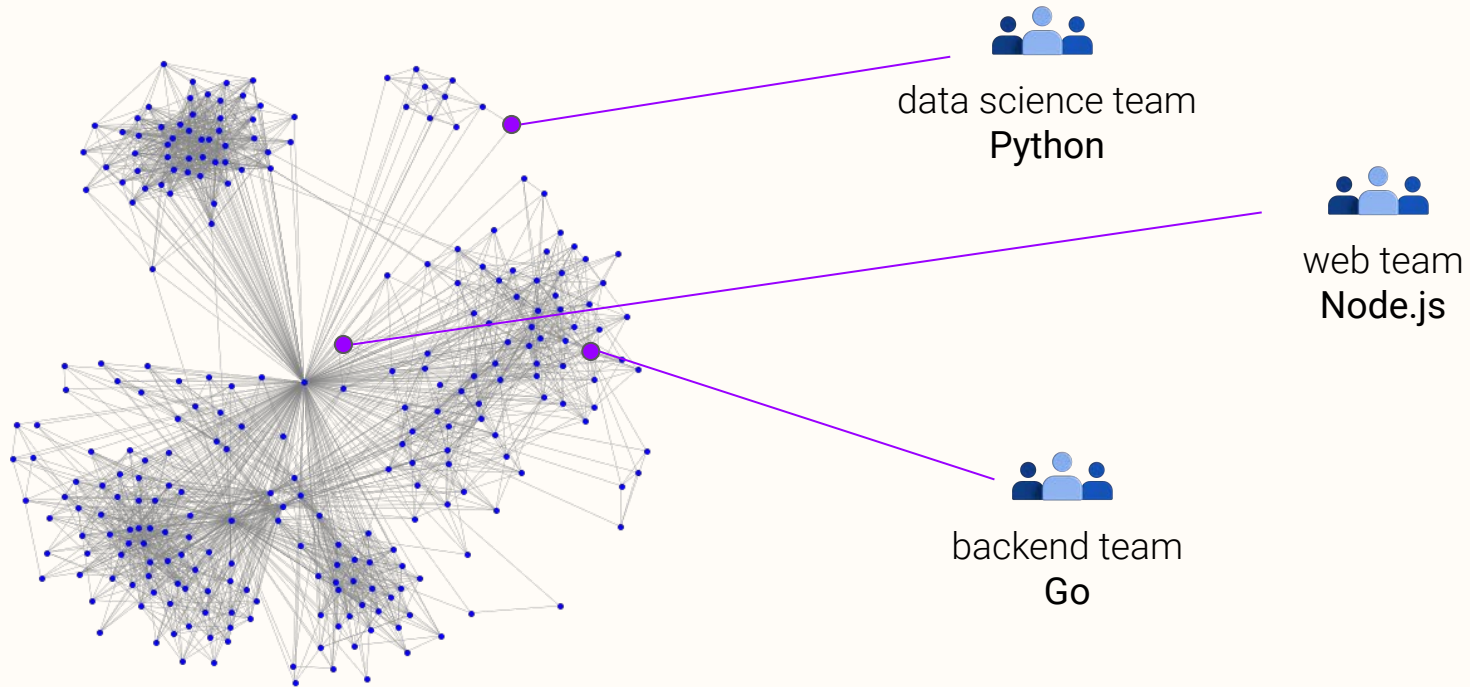
dev integrate solution



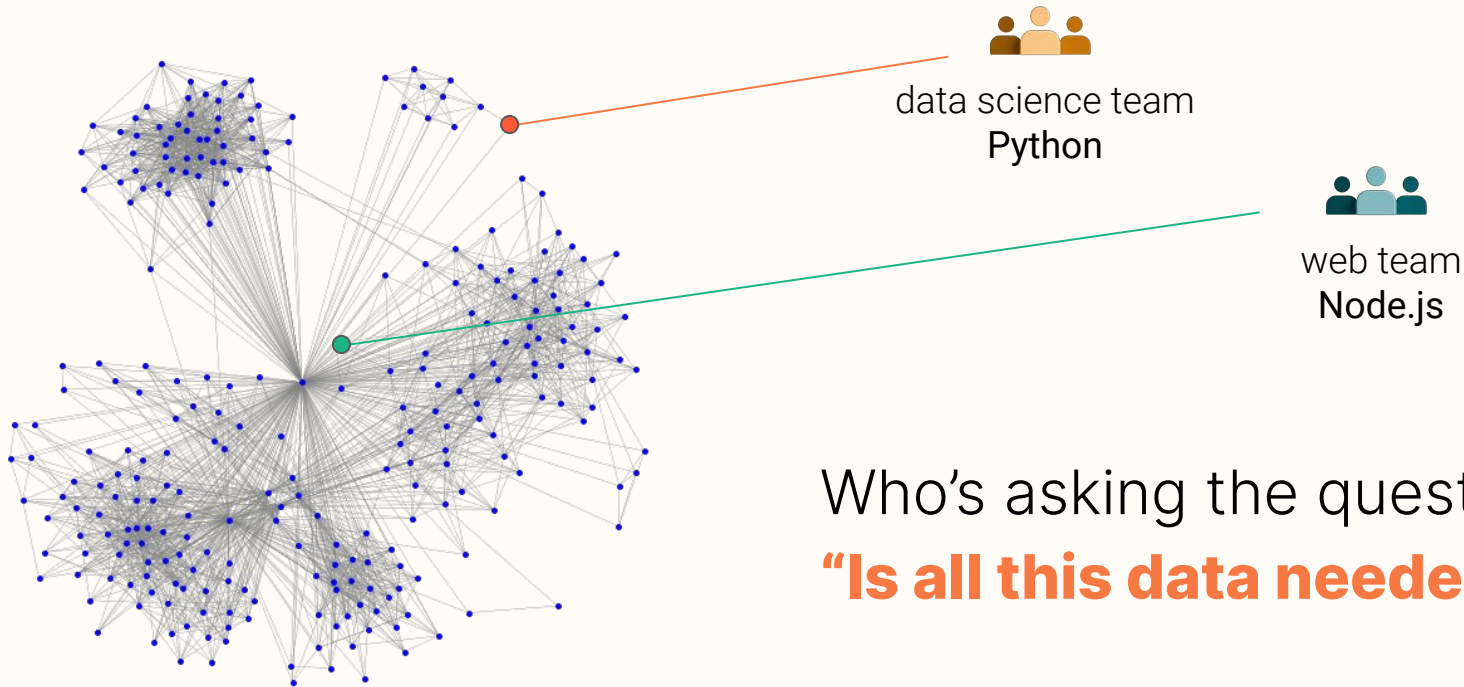
dev decide what to “measure”




dev build dashboards etc.








Who's asking the question  
**"Is all this data needed?"**



Worried DevOps / SRE



Are we giving who's responsible  
**the Tools to Succeed?**

→ get 100% coverage?


→ ensure cost - visibility trade-off

▼ Unscalable **Pricing** Models

▼ Harder Org **Alignment**

▼ Data **Privacy**

**Legacy Observability  
for Cloud-Native**



## 011y is Under-Adopted

Devops Pulse, 2022

# That's why we've built **groundcover**.



Data is collected  
**out-of-band with eBPF**



Data is digested  
**distributedly where it lies**



In-cloud architecture where  
**All data is kept private**



## **O11y is hard to integrate.**

Requiring lots of R&D effort around code instrumentations, and complex coordination between teams


## Auto-instrumentation...

```
// Run starts polling users for Fibonacci number requests and writes results.
func (a *App) Run(ctx context.Context) error {
    for {
        // Each execution of the run loop, we should get a new "root" span and context.
        newCtx, span := otel.Tracer(name).Start(ctx, "Run")


        n, err := a.Poll(newCtx)
        if err != nil {
            span.End()
            return err
        }

        a.Write(newCtx, n)
        span.End()
    }
}
```

# Back to the **the dev cycle**



**Instrument code...**  
Web team (Node.js)

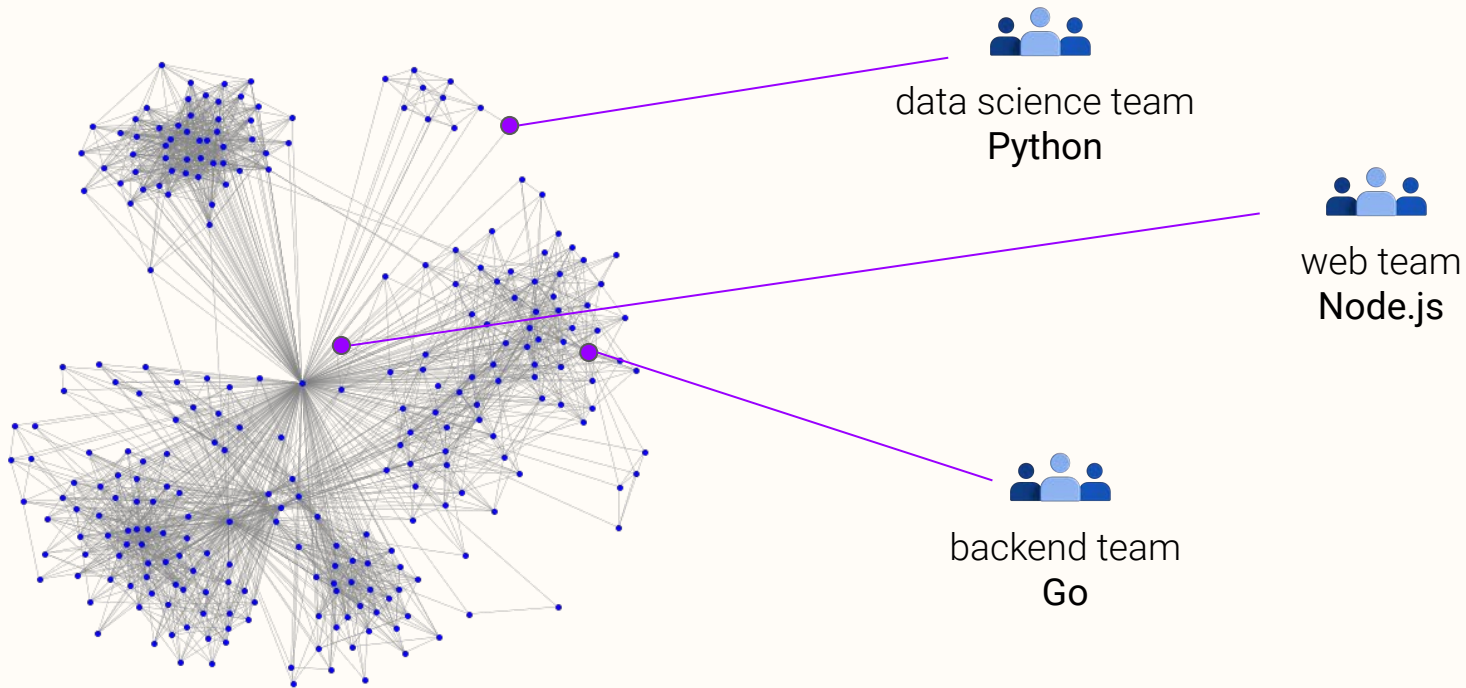


**Wait for version release...**




**Get value in production (!)**






Now try this in a **real company**



Worried DevOps / SRE



Are we giving who's responsible  
**for aligning all these teams?**

→ to one uniform approach

→ to a high professional standard



**eBPF does to Linux what JavaScript does to HTML...with eBPF, instead of a fixed kernel, you can now write mini programs...which are run in a safe virtual machine in the kernel."**

Brendan Gregg, production eng. @Netflix

## **eBPF, the next Linux superpower**

# eBPF in 60 secs

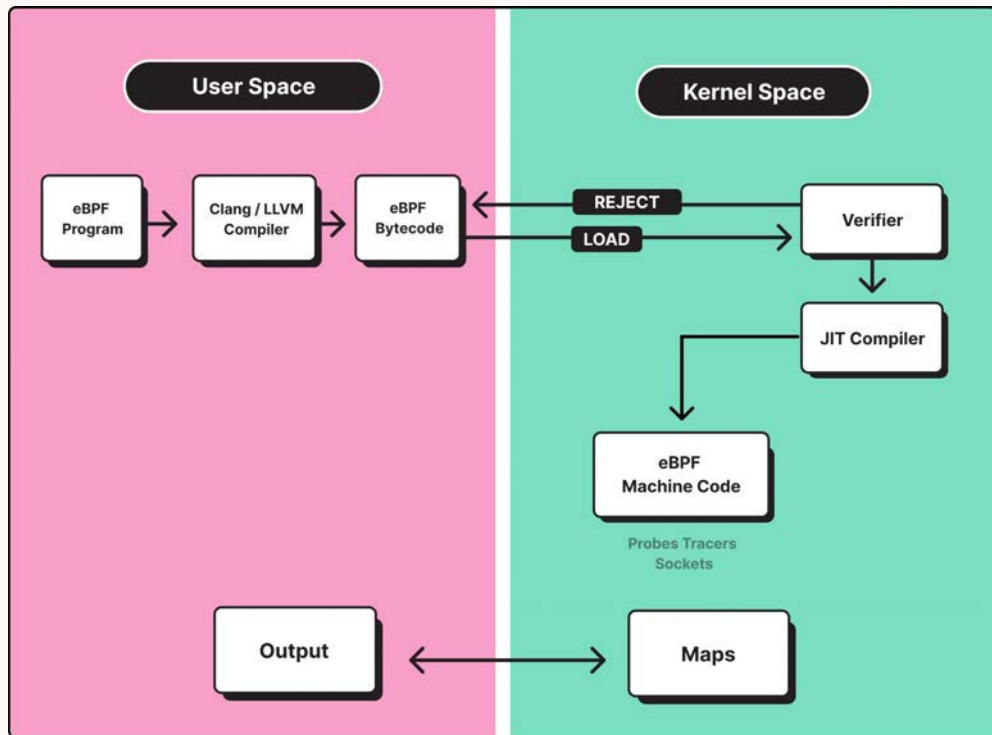


● BPF | 1993

● eBPF | 2014

● v4.14 | 2017

# eBPF architecture



## ▲ **Efficiency**

Ever thought about your O11y SDK overhead?


## ▲ **Safety**

You can't crash the application.


## ▲ **100% coverage**

You can see all user-space apps, out-of-band

**But why run things  
in the kernel anyway?**



Happy DevOps / SRE



No more convincing **the R&D**

- one uniform approach
- super high professional standard

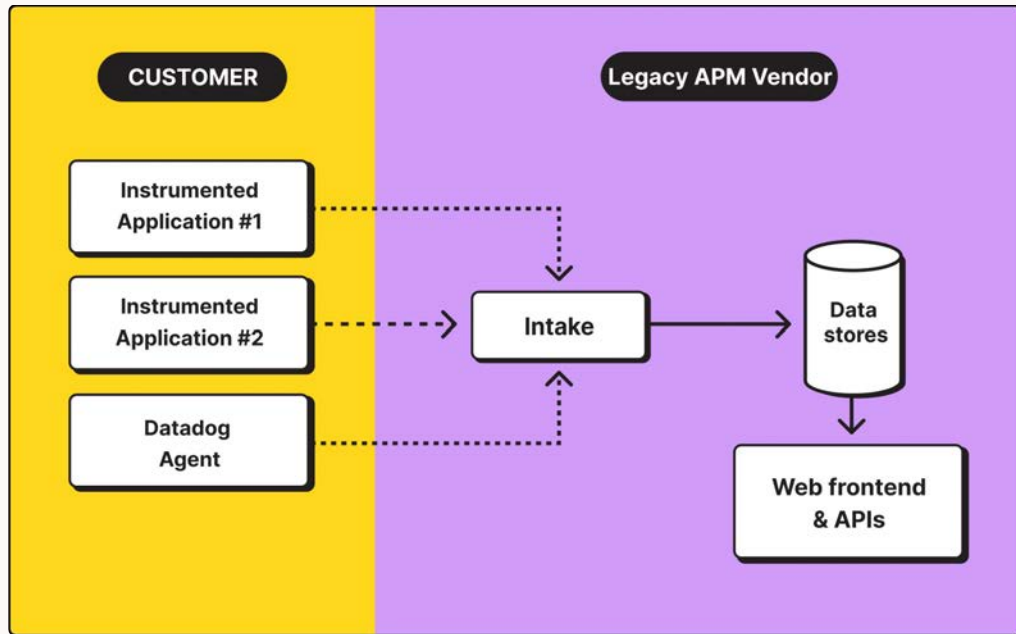


## **011y doesn't scale.**

Storing huge amounts of irrelevant data for the little insight you need, making costs unbearable



# Centralized architecture



## ▼ **Random Sampling**

“How can I control data volumes?”

I have high-throughput APIs”

## ▼ **Raw data is stored**

“How do I get a P50 latency metric over time?”

Span-based metric at the vendor’s backend

## ▼ **Rigid data collection**

“OK this is an interesting! give me everything!”

So you mean you want it all the time?

**The debt of the  
centralized architecture**



# Limiting your cardinality


Where you do care.

Status Code	Latency
200 - OK	29 ms
200 - OK	93.6 ms
200 - OK	7.8 ms
200 - OK	6.8 ms
200 - OK	228 $\mu$ s
200 - OK	292 $\mu$ s
200 - OK	103 $\mu$ s
200 - OK	63.2 ms
200 - OK	189 $\mu$ s
404 - Not Found	217 $\mu$ s
404 - Not Found	217 $\mu$ s
200 - OK	161 $\mu$ s
200 - OK	98.2 $\mu$ s
200 - OK	43.6 ms
200 - OK	669 $\mu$ s
200 - OK	4.3 ms
200 - OK	98.5 ms
200 - OK	40.5 ms
200 - OK	75 $\mu$ s
200 - OK	66.4 ms
200 - OK	9.6 ms
404 - Not Found	243 $\mu$ s
200 - OK	25.9 ms
200 - OK	5.4 ms






## In-host span-based metrics



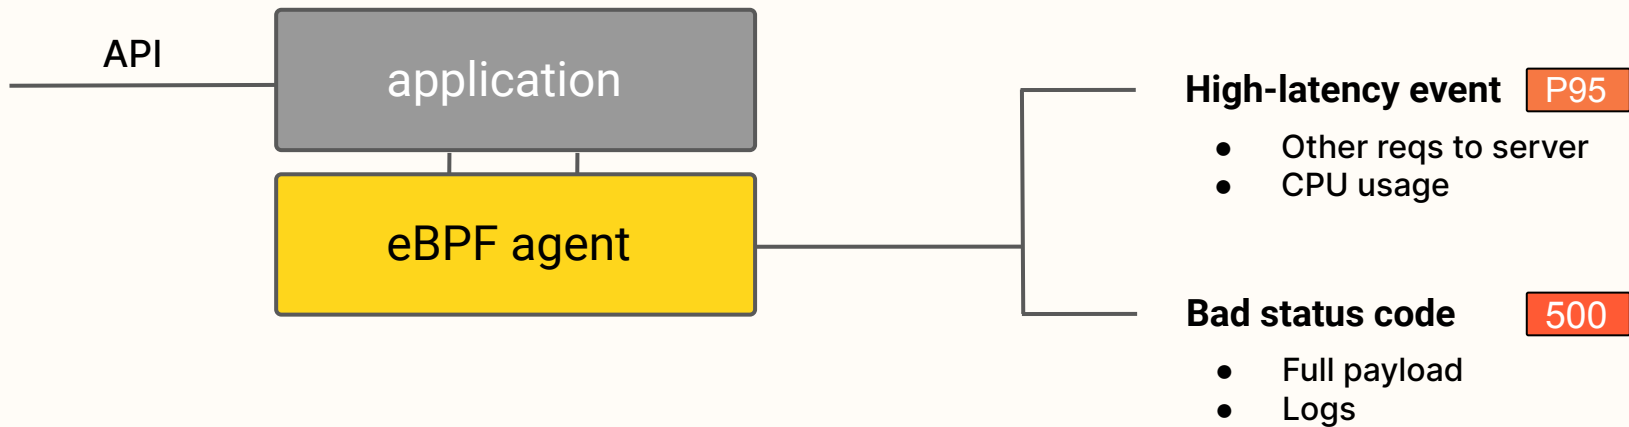


## Variant-depth capturing



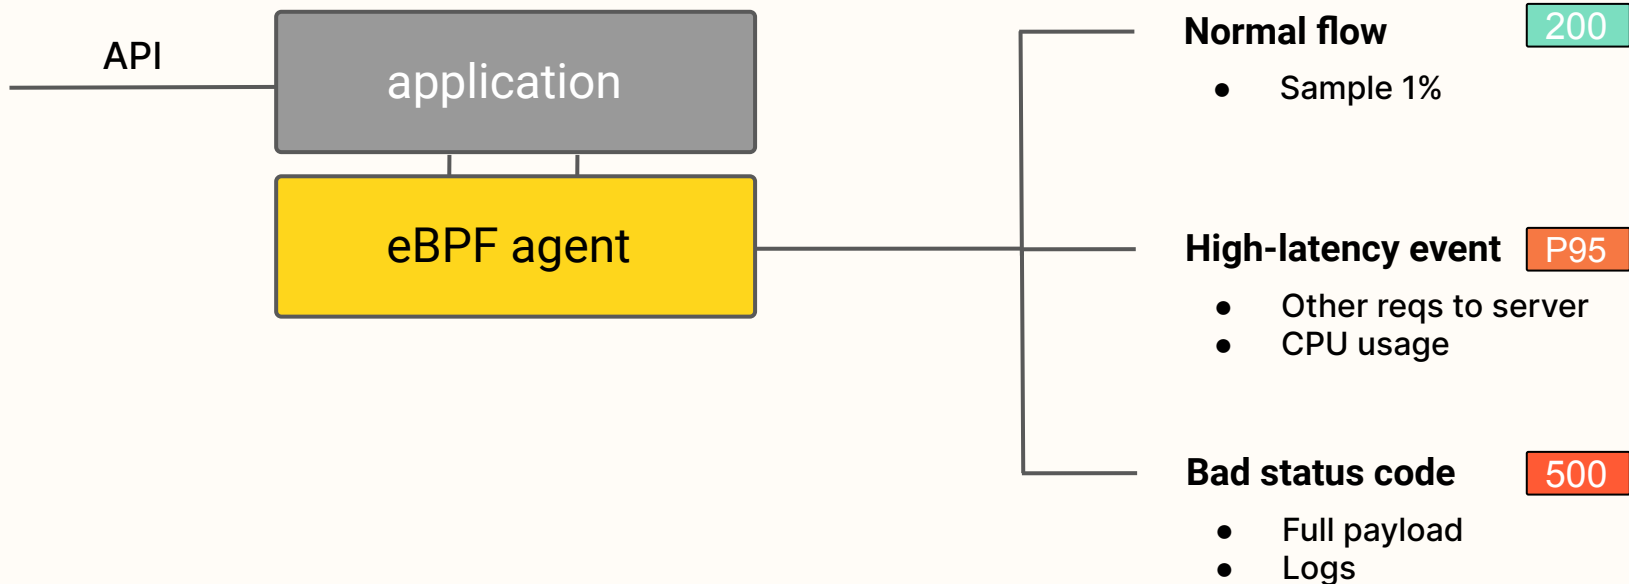


## Variant-depth capturing





## Variant-depth capturing

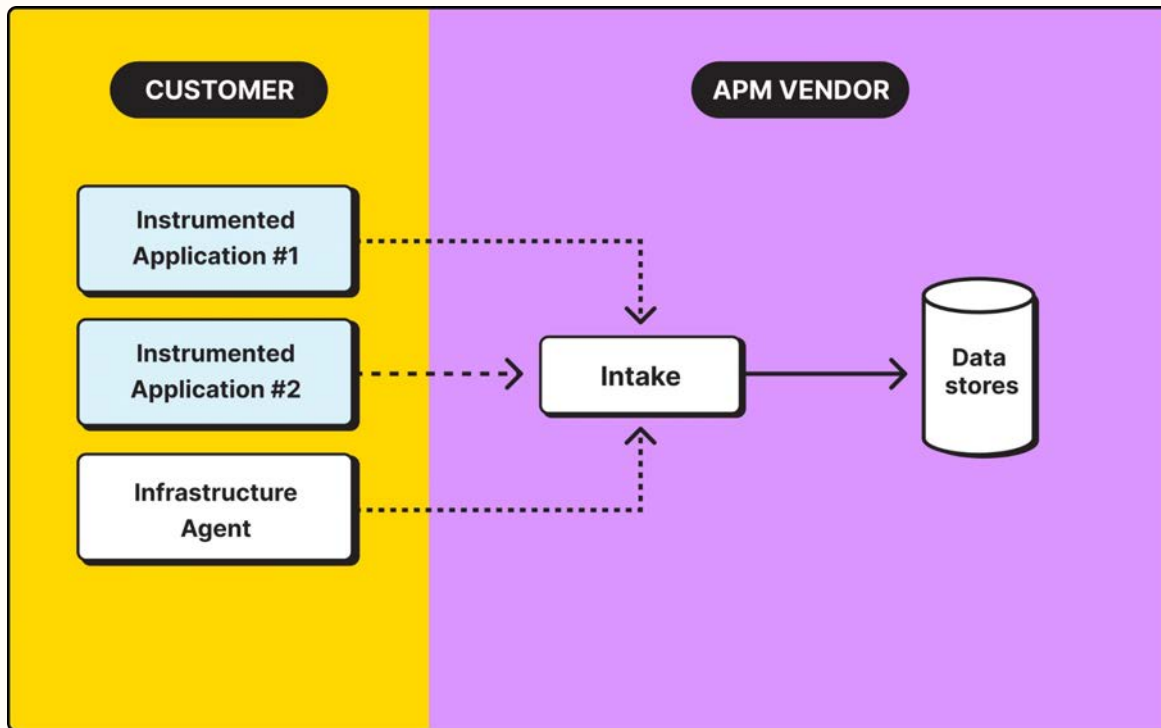


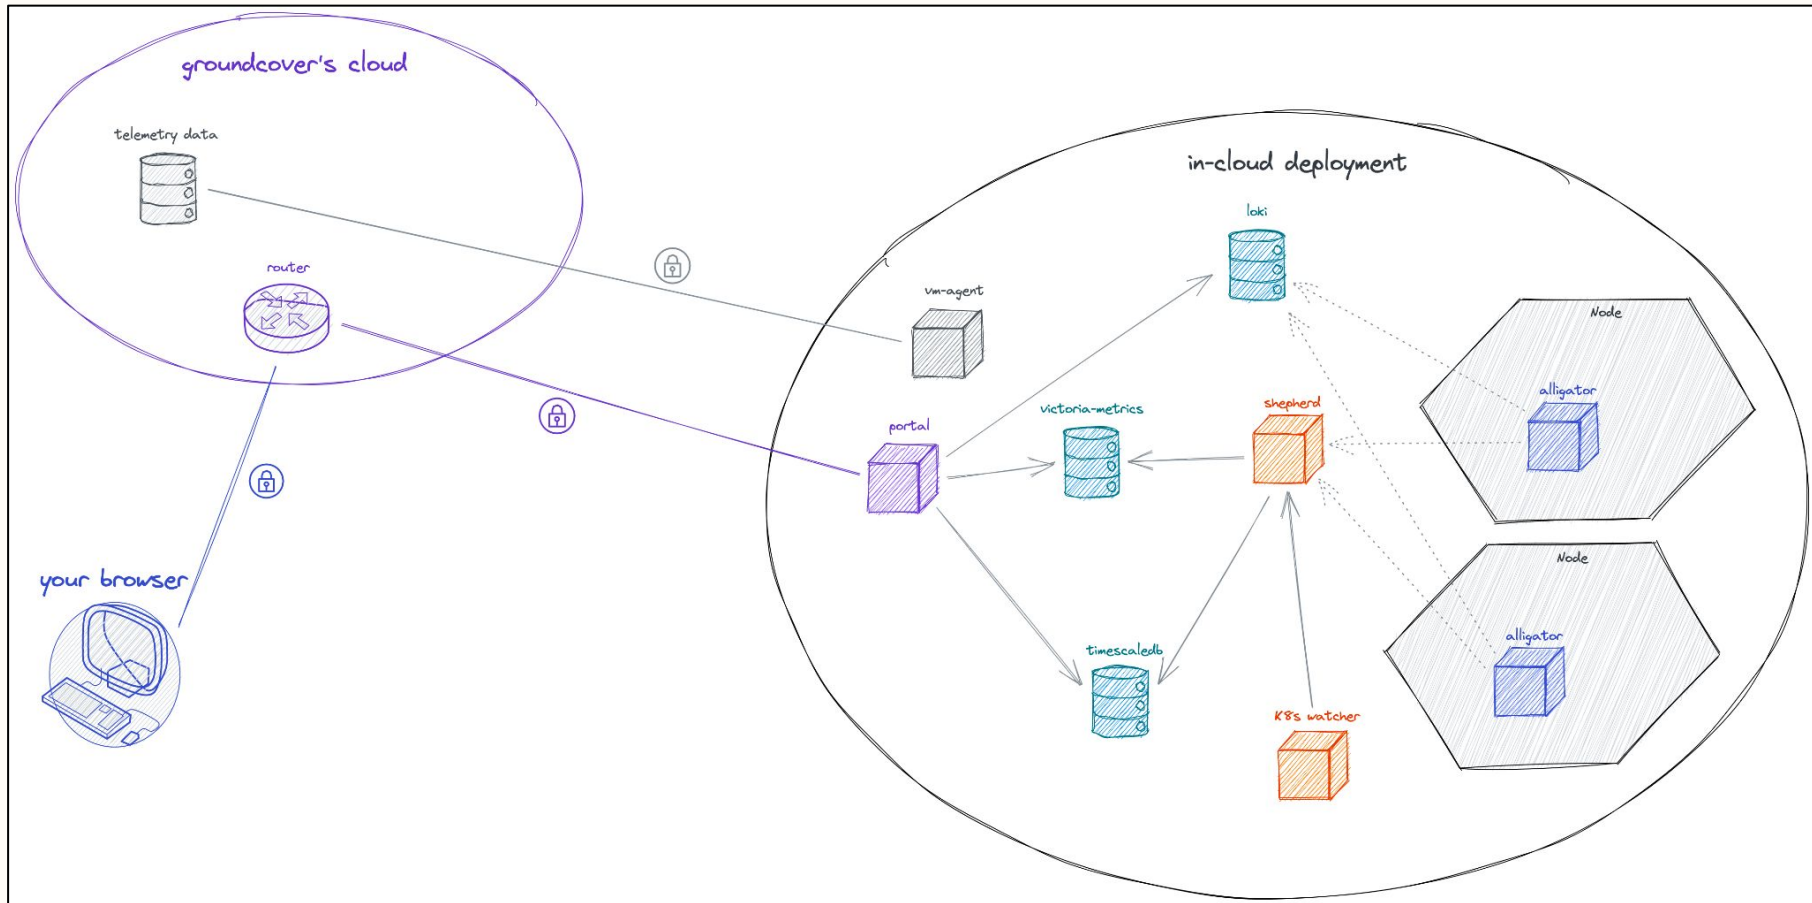




## **Your data is not kept private.**

Storing sensitive information like logs and full traces at the vendor's side





# This is our current reality.

- **eBPF instrumentation**
  - Immediate time-to-value
  - Out-of-band deployment
- **Edge-based observability**
  - Built for scale
  - Breaks all tradeoffs
- **In-cloud architecture**
  - Data is kept private
  - In your full control



groundcover



# Start Free



[www.groundcover.com](http://www.groundcover.com)

