entrinsicon

# Informer Cloud

## And Multi-Tenant Informer

**Derek O'Neill**

Lead DevOps Engineer

# Agenda

entrinsicon

1. Who?
2. How?
3. Why?
4. Why Not?
5. Huh? (Q&A)

# Who?

- Derek O'Neill
- Lead DevOps Engineer at Entrinsik: 5 years
- Software Engineer on Informer 5: 3 years
- Informer Implementation Specialist: 2 years

# How?


Glossary:

- Cloud – Someone else's computer(s)
- Kubernetes – Software to manage containers and infrastructure at scale
- Containers – Small almost-VMs that run single pieces of software
- Docker – A software system for building and deploying containers

# High-Level Infrastructure

entrinsicon

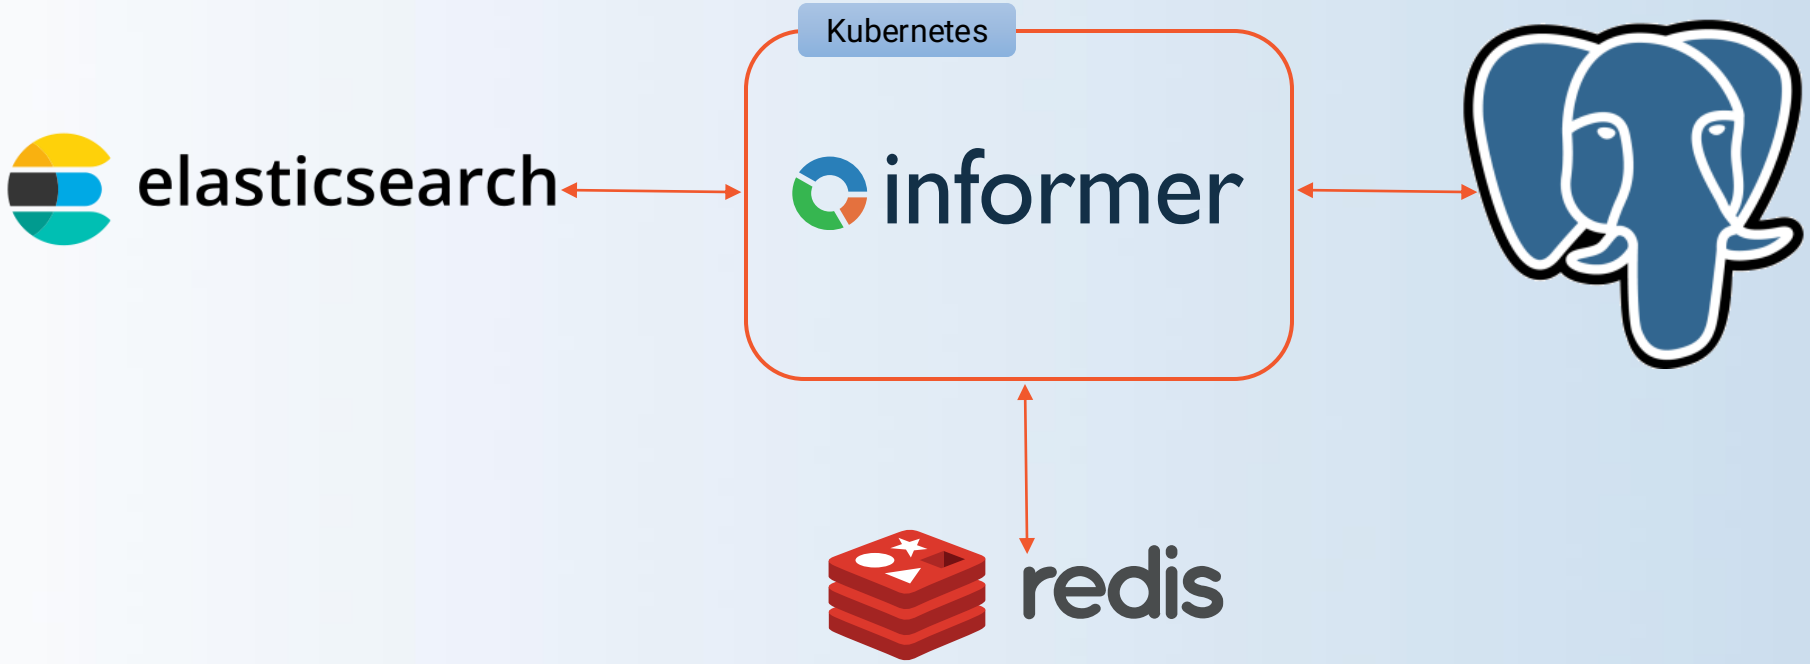

# Siloing

- Having dedicated stacks for each customer is not scalable.
- However, data and metadata siloing are *essential*.
- Resolve the tension with clusterId and leverage it in conjunction with PG, ES, and Redis' built-in security features.

# Resource Contention

This is the hard part, and every release requires some degree of re-engineering. But here are the highlights of what we've learned.

1. Jobs should be run in separate processes with their own pool of resources.
2. Queries might need to be as well depending on usage patterns.
3. There's no such thing as too much memory at scales like Informer Cloud.
4. Tuning liveness, readiness, and scaling parameters is an ever-evolving process.
5. Kubernetes is not a silver bullet. It's a tool that needs to be used by a skilled engineer with specific goals.

# Why?

- Economics
- Compliance
- Performance
- Uptime
- Seamless upgrades

# Why Not?


- Very specific compliance situations
- Custom architecture requirements
- Unfamiliar bottlenecks

# Huh? (Q&A)

Ask me anything.