Building a Cloud-Native Technology Stack that Supports Full Cycle Development

Daniel Bryant

Product Architect, Ambassador Labs (formerly Datawire)

DATAWIRE

tl;dr

- Being fully cloud native requires new tech and new workflows

- Creating a supporting cloud platform is essential:
  - Container orchestration
  - Progressive delivery
  - Edge management
  - Observability

- Consciously design your platform & watch for antipatterns

@danielbryantuk

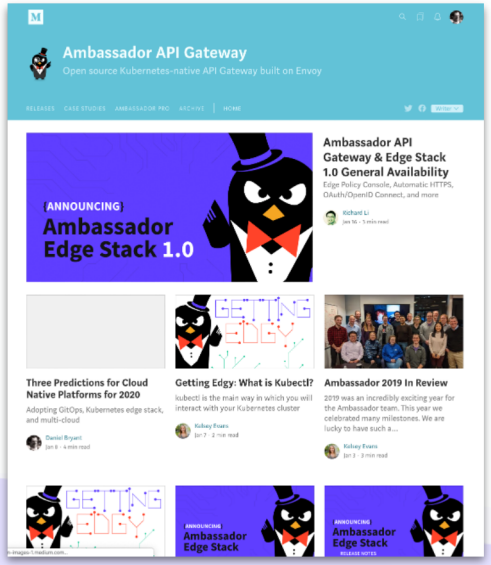

A quick cloud native primer...

- Going “cloud native” offers benefits, but requires changes:
  - New technologies
  - Appropriate culture
  - New workflows

- Successful cloud native organisations have:
  - Created a self-service application platform
  - Adopted new tools and (full cycle) developer workflows

Full Cycle Developers

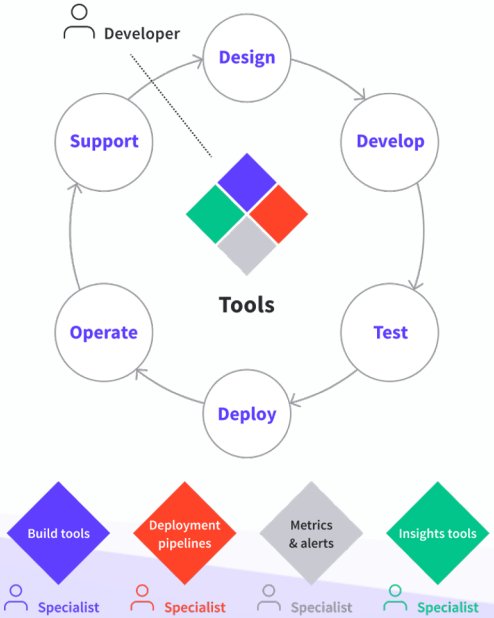

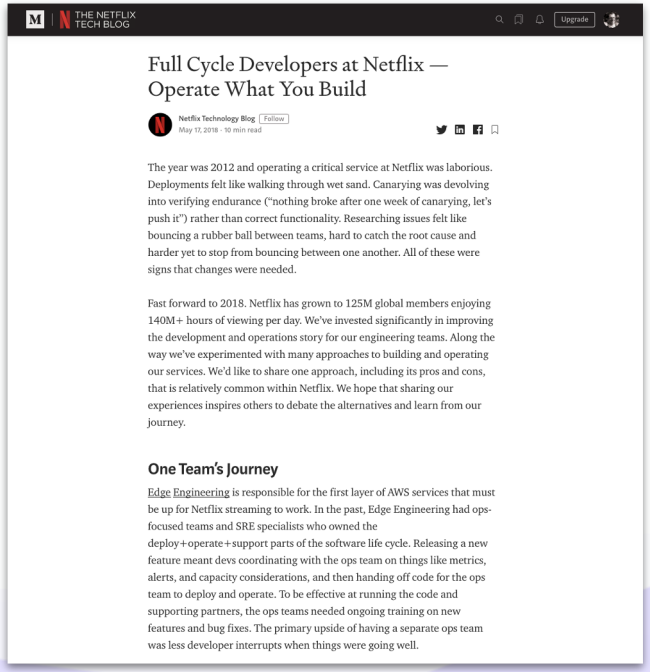

THE NETFLIX TECH BLOG

Upgrade

Full Cycle Developers at Netflix — Operate What You Build

Netflix Technology Blog · Follow
May 17, 2018 · 10 min read

The year was 2012 and operating a critical service at Netflix was laborious. Deployments felt like walking through wet sand. Canarying was devolving into verifying endurance ("nothing broke after one week of canarying, let's push it") rather than correct functionality. Researching issues felt like bouncing a rubber ball between teams, hard to catch the root cause and harder yet to stop from bouncing between one another. All of these were signs that changes were needed.

Fast forward to 2018. Netflix has grown to 125M global members enjoying 140M+ hours of viewing per day. We've invested significantly in improving the development and operations story for our engineering teams. Along the way we've experimented with many approaches to building and operating our services. We'd like to share one approach, including its pros and cons, that is relatively common within Netflix. We hope that sharing our experiences inspires others to debate the alternatives and learn from our journey.

One Team's Journey

Edge Engineering is responsible for the first layer of AWS services that must be up for Netflix streaming to work. In the past, Edge Engineering had ops-focused teams and SRE specialists who owned the deploy+operate+support parts of the software life cycle. Releasing a new feature meant devs coordinating with the ops team on things like metrics, alerts, and capacity considerations, and then handing off code for the ops team to deploy and operate. To be effective at running the code and supporting partners, the ops teams needed ongoing training on new features and bug fixes. The primary upside of having a separate ops team was less developer interrupts when things were going well.

https://netflixtechblog.com/full-cycle-developers-at-netflix-a08c31f83249

Full Cycle Developers: Team Topologies

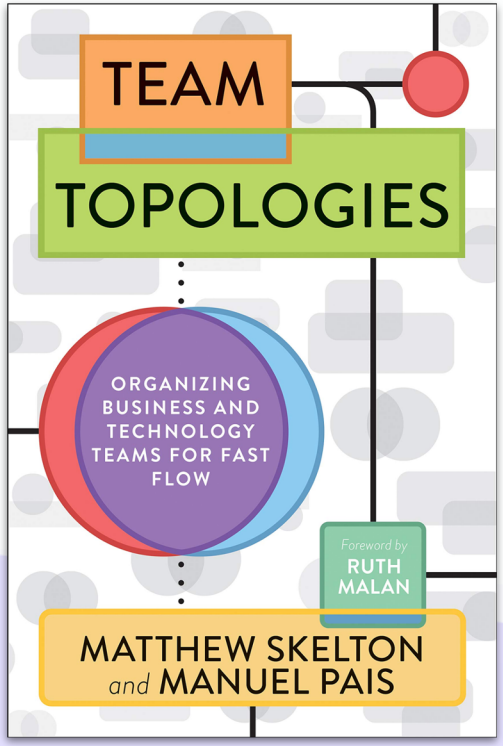

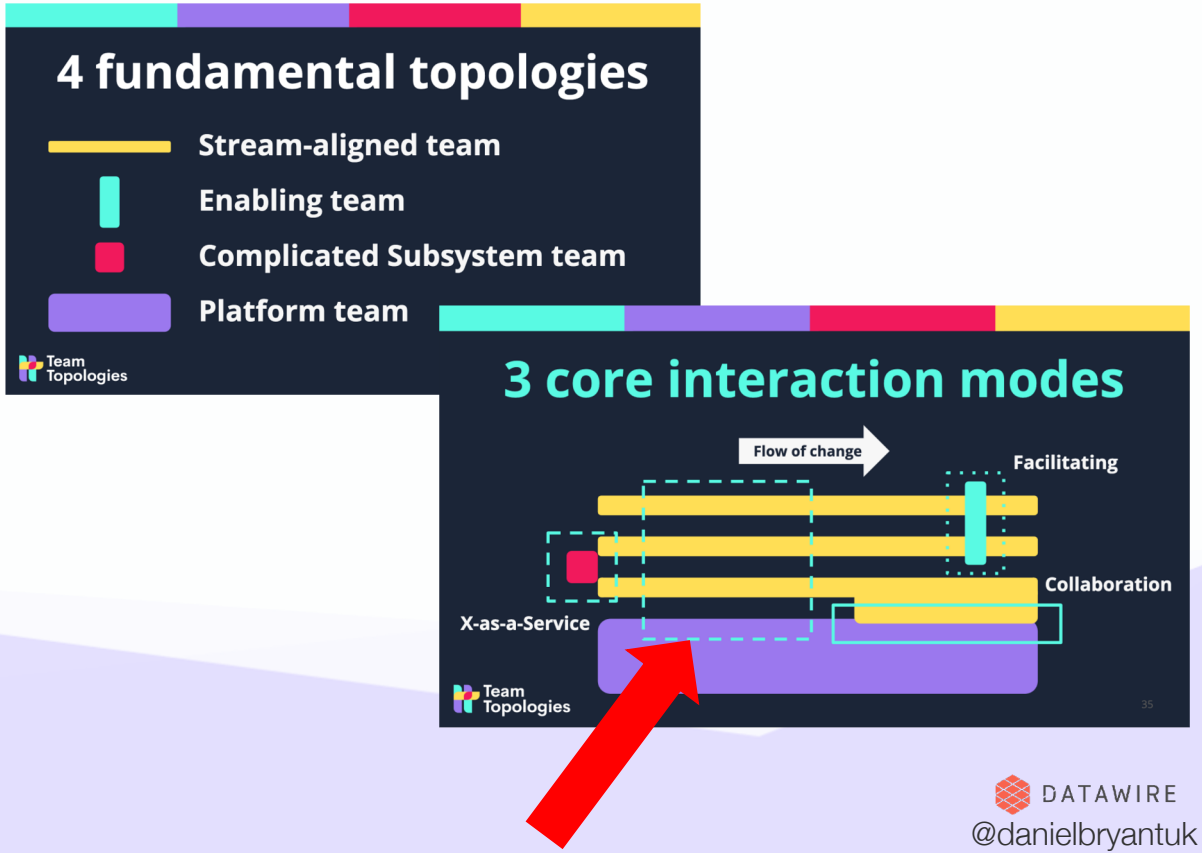

DATAWIRE
@danielbryantuk

Four cloud native platform requirements

1. Container management

1. Progressive delivery

1. Edge management

1. Observability

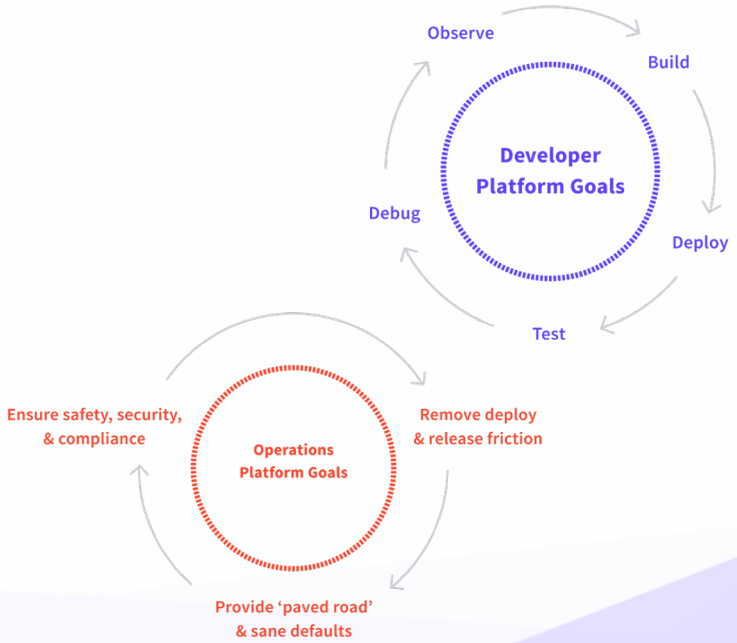

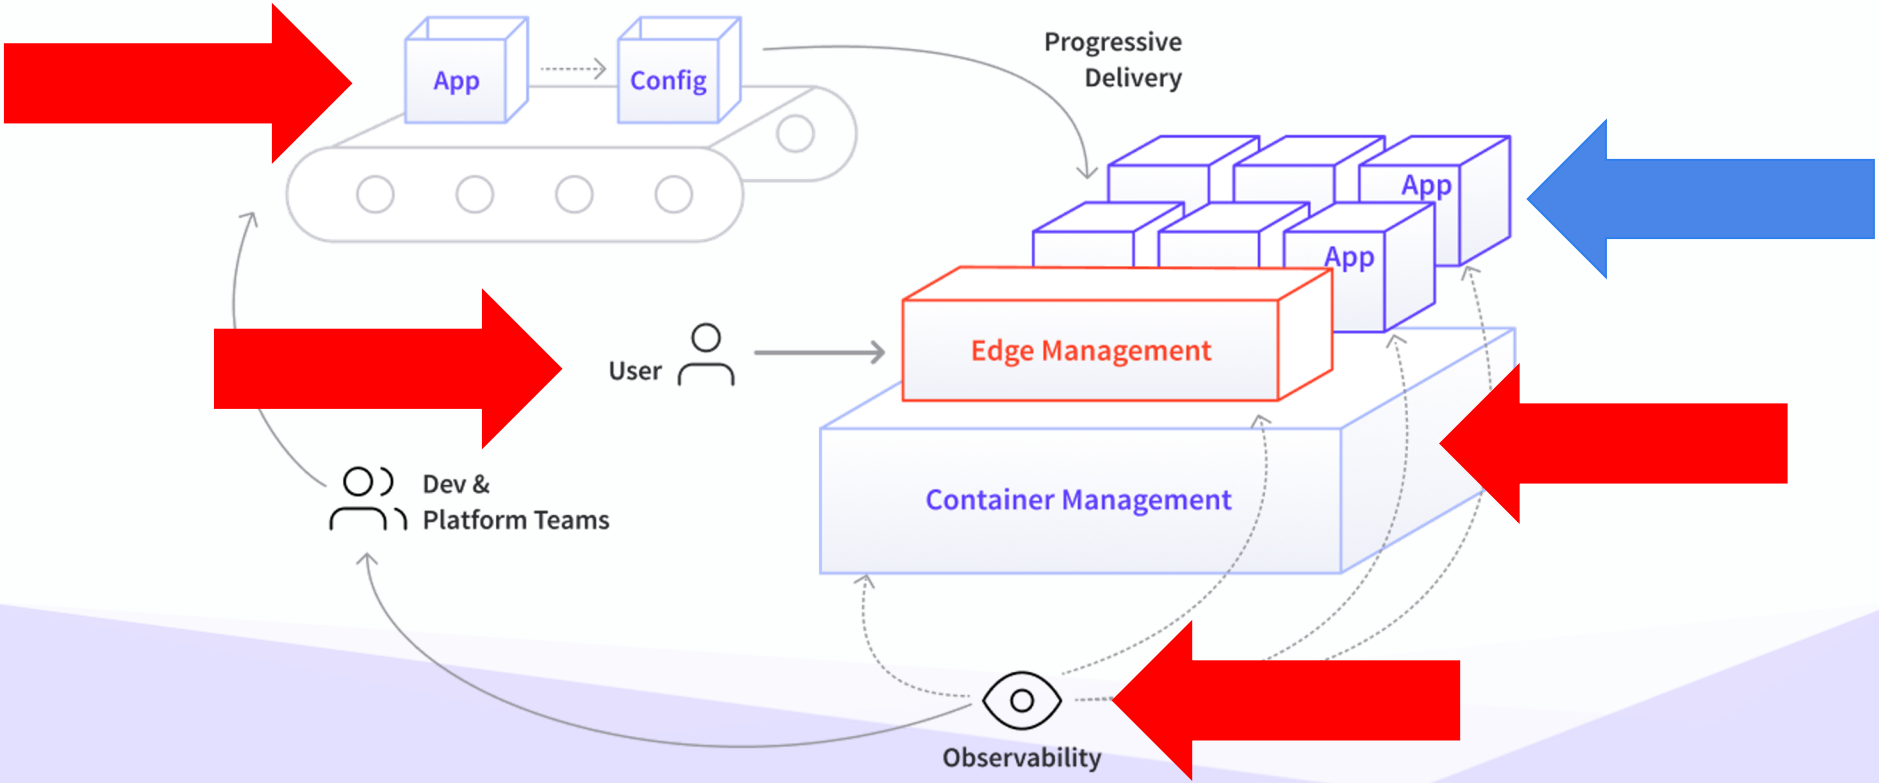
App
Config
Progressive
Delivery
App
App
User
Edge Management
Container Management
Dev &
Platform Teams
Observability

More Details on Full Cycle and K8s

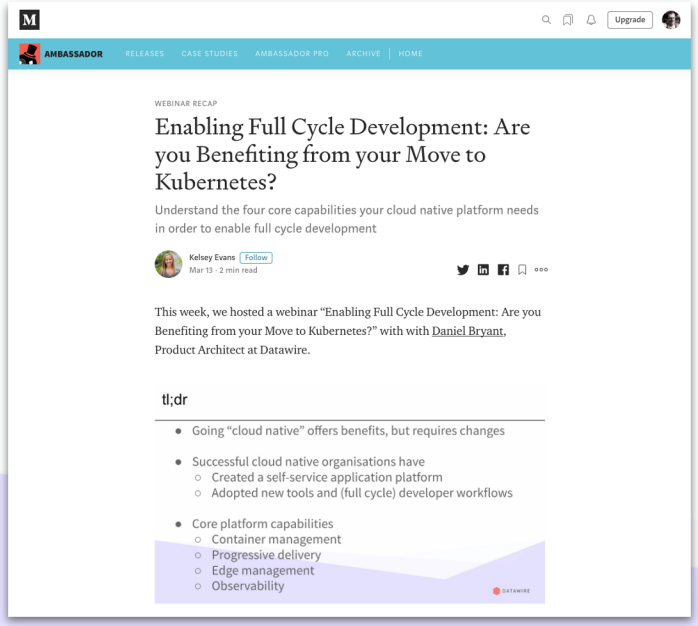

- Successful cloud native organisations have:
  - Created a self-service application platform
  - Adopted new tools and (full cycle) developer workflows

https://blog.getambassador.io/enabling-full-cycle-development-are-you-benefiting-from-your-move-to-kubernetes-d9eab2e94e7

Avoiding
Platform Antipatterns

DATAWIRE

Avoiding Platform Antipatterns

Centralized Control and Ownership: One Size Doesn’t Fit All

Fragmented Platform Implementation

Slow Development Loops: Less Time Coding, More Time Toiling

Antipattern: Centralized Control and Ownership

- (Dis)economies of scale
- Overzealous guardrails
- Modification is ticket-driven

Antipattern: Fragmented Platform Implementation

Pattern: Microservices

Description

Design modules as separate deployment and operation units, with large degrees of freedom for their implementation

Approach

- Former technical detail (deployment architecture) as first class architectural design principle
- Network communication as hard-to-cross boundary, enforcing encapsulation

Consequences

- Isolation
- Autonomy
- Scalability
- Resilience
- Speed
- Experimentation
- Rapid Feedback
- Flexibility
- Replaceability

Antipattern: Micro Platform

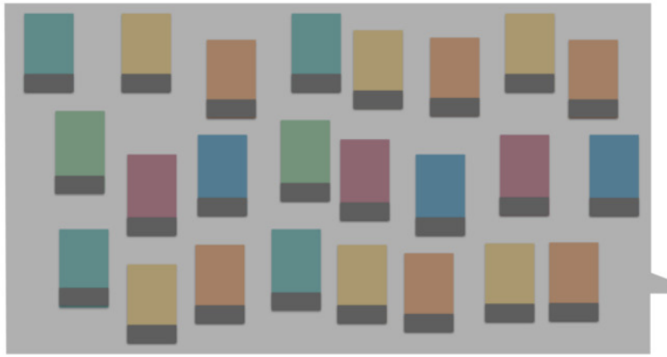

https://speakerdeck.com/stilkov/microservices-patterns-and-antipatterns-1?slide=12

DATAWIRE

Antipattern: Slow Development Loops

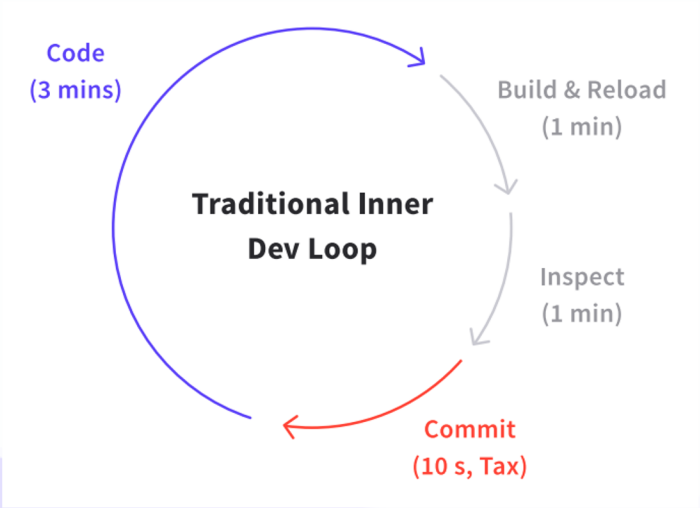

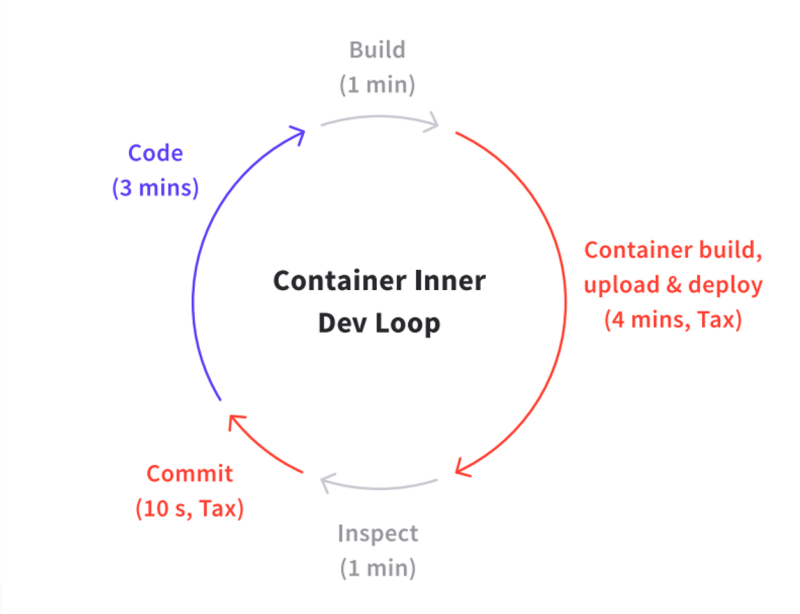

https://mitchdenny.com/the-inner-loop/

Exploring the Platform Capabilities

DATAWIRE

Four Core Platform Capabilities

1. Container management

1. Progressive delivery

1. Edge management

1. Observability

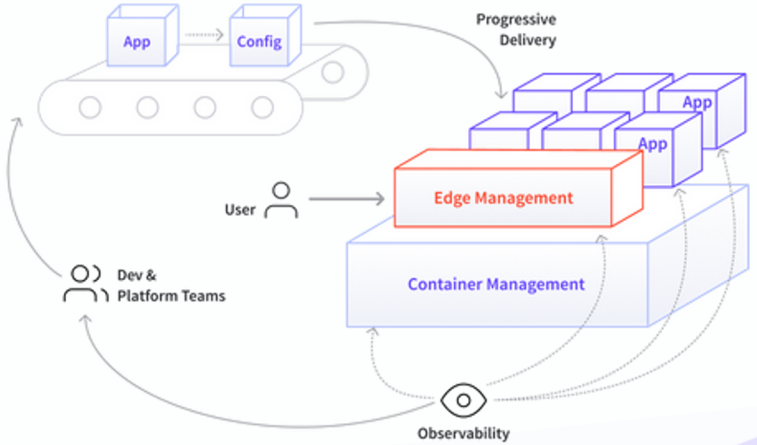

Container Management: Kubernetes

DATAWIRE

Container Management

Manage and run container-based applications at scale and on a variety of infrastructures

- Developers
  - Self-service interactions: automated and observable

- Platform team
  - Set policies around access, control, and auditability

Kubernetes Decisions

- To self-host, or not to self-host?

- Which distro?

- Going all-in on a cloud?

Google Kubernetes Engine

Azure Kubernetes Service (AKS)

DATAWIRE

@danielbryantuk

Kubernetes Challenges

- Foundations for a PaaS-like experience?
  - Helm and Helmfile for deployment

- Developer productivity
  - Local-to-remote dev and test

Progressive Delivery: Delivery Pipelines

Progressive Delivery

Supporting the creation of pipelines that enable the automated build, verification, deployment, release, and observability

- Developers
  - Self-service interactions: automated and observable

- Platform team
  - Centralize verification of quality and security properties

https://redmonk.com/jgovernor/2018/08/06/towards-progressive-delivery/

Progressive Delivery Decisions

- Deliver any and all application changes into production as **rapidly** and as **safely** as the organisation requires
  - Pipeline practices
  - Pipeline technology

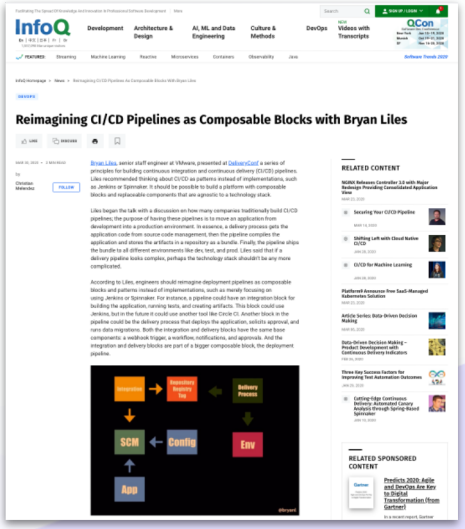

https://www.infoq.com/news/2020/03/reimagining-cicd-pipelines/

Progressive Delivery Challenges

- Collaboration between dev, QA, and ops
- Balance one-size-fits-all vs chaos
- Make it easy to do the right thing

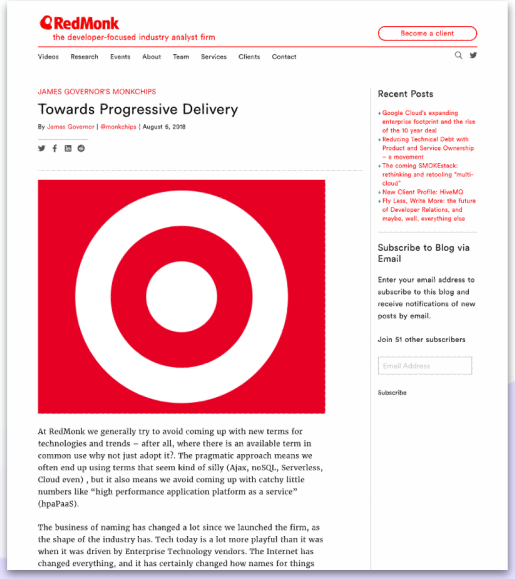

Edge Management: Ingress and API Gateways

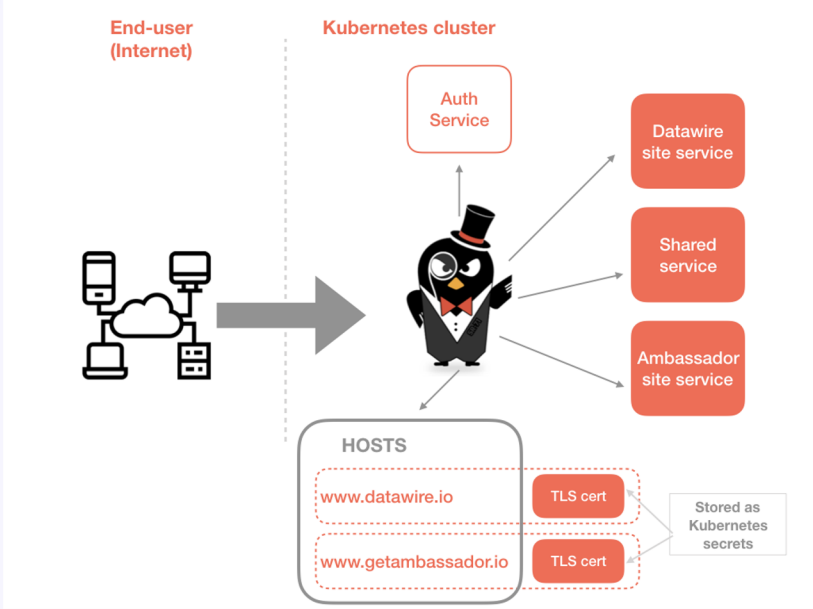

DATAWIRE

Edge Management

Enable the self-service release of new functionality by developers, while maintaining stability

- Developers
  - Decentralized traffic management
  - Support NFRs e.g. authn/z, retries, and circuit breaking

- Platform
  - Centralized configuration of sane defaults
  - TLS, authn/z, and rate limiting for DDoS protection

Edge Stack Decisions

- Edge technologies
  - Envoy becoming the de facto standard(?)
  - xDS APIs / Ingress v2

- Deploy/release workflows
  - Declarative (CRDs)
  - Self-service

Edge Stack Challenges

- Scaling edge management
- Supporting multiple protocols and NFRs

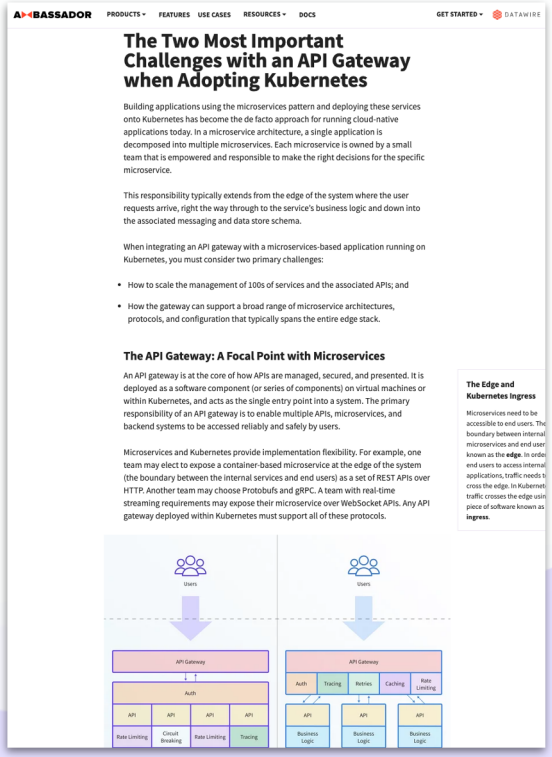

AMBASSADOR PRODUCTS FEATURES USE CASES RESOURCES DOCS GET STARTED DATAWIRE

The Two Most Important Challenges with an API Gateway when Adopting Kubernetes

Building applications using the microservices pattern and deploying these services onto Kubernetes has become the de facto approach for running cloud-native applications today. In a microservice architecture, a single application is decomposed into multiple microservices. Each microservice is owned by a small team that is empowered and responsible to make the right decisions for the specific microservice.

This responsibility typically extends from the edge of the system where the user requests arrive, right the way through to the service's business logic and down into the associated messaging and data store schema.

When integrating an API gateway with a microservices-based application running on Kubernetes, you must consider two primary challenges:

- How to scale the management of 100s of services and the associated APIs; and
- How the gateway can support a broad range of microservice architectures, protocols, and configuration that typically spans the entire edge stack.

The API Gateway: A Focal Point with Microservices

An API gateway is at the core of how APIs are managed, secured, and presented. It is deployed as a software component (or series of components) on virtual machines or within Kubernetes, and acts as the single entry point into a system. The primary responsibility of an API gateway is to enable multiple APIs, microservices, and backend systems to be accessed reliably and safely by users.

Microservices and Kubernetes provide implementation flexibility. For example, one team may elect to expose a container-based microservice at the edge of the system (the boundary between the internal services and end users) as a set of REST APIs over HTTP. Another team may choose Protobufs and gRPC. A team with real-time streaming requirements may expose their microservice over WebSocket APIs. Any API gateway deployed within Kubernetes must support all of these protocols.

The Edge and Kubernetes Ingress

Microservices need to be accessible to end users. The boundary between internal microservices and end users known as the **edge**. In order end users to access internal applications, traffic needs t cross the edge. In Kubernet traffic crosses the edge usin piece of software known as **ingress**.

Users Users

API Gateway API Gateway

Auth Auth Tracing Retries Caching Rate Limiting

API API API API API API API API

Rate Limiting Circuit Breaking Rate Limiting Tracing Business Logic Business Logic Business Logic

https://www.getambassador.io/resources/challenges-api-gateway-kubernetes/

DATAWIRE

@danielbryantuk

Observability:
Metrics, Logging, Tracing

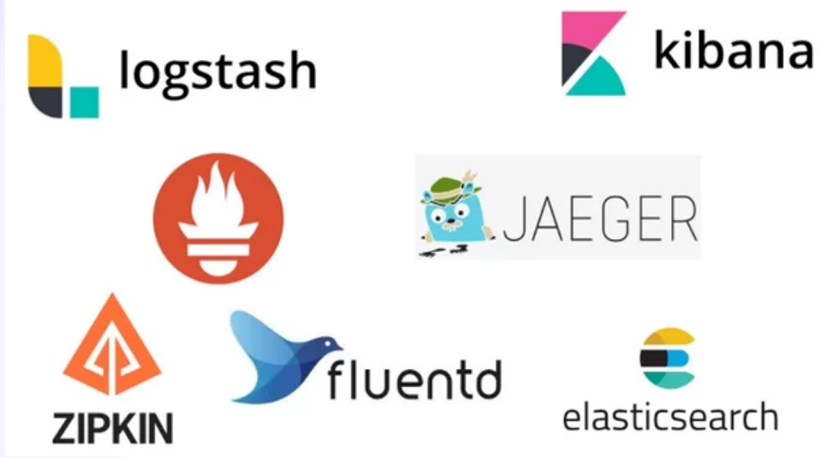

Observability

Support the collection and analysis of end user and application feedback directly by developers and the platform team.

- Developers
  - Enable product teams to observe and iterate against business goals and KPIs

- Platform
  - Observe and managing infrastructure, and ensure their service level objectives (SLOs) are met

Observability Decisions

- Adoption (monitor all-the-things?)
- Technology selection (standards)
  - Metrics
  - Logging
  - Distributed tracing
- Joining the dots

LIGHTSTEP

Observability Challenges

- Self-service config and dashboards
- Increasing signal-to-noise
- Fault location

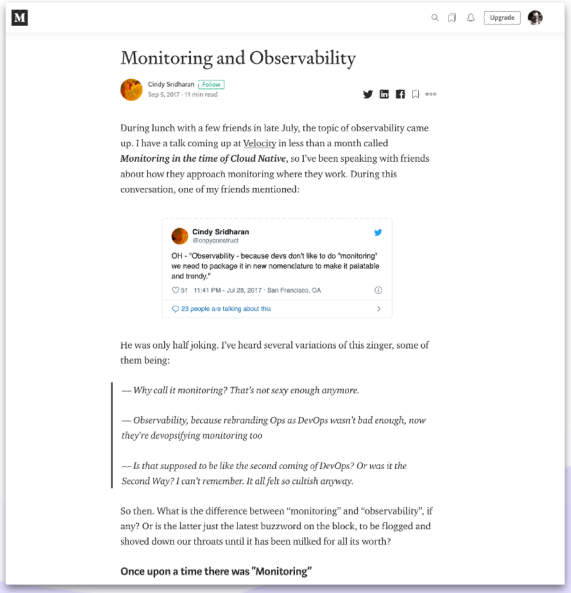

Monitoring and Observability

Cindy Sridharan · Follow · Sep 5, 2017 · 11 min read

During lunch with a few friends in late July, the topic of observability came up. I have a talk coming up at Velocity in less than a month called ***Monitoring in the time of Cloud Native***, so I've been speaking with friends about how they approach monitoring where they work. During this conversation, one of my friends mentioned:

> **Cindy Sridharan** @copyconstruct
>
> Oh - "Observability - because devs don't like to do "monitoring" we need to package it in new nomenclature to make it palatable and trendy."
>
> 51 11:41 PM - Jul 28, 2017 - San Francisco, CA
>
> 23 people are talking about this

He was only half joking. I've heard several variations of this zinger, some of them being:

> *— Why call it monitoring? That's not sexy enough anymore.*
>
> *— Observability, because rebranding Ops as DevOps wasn't bad enough, now they're devopsifying monitoring too*
>
> *— Is that supposed to be like the second coming of DevOps? Or was it the Second Way? I can't remember. It all felt so cultish anyway.*

So then. What is the difference between "monitoring" and "observability", if any? Or is the latter just the latest buzzword on the block, to be flogged and shoved down our throats until it has been milked for all its worth?

Once upon a time there was "Monitoring"

https://medium.com/@copyconstruct/monitoring-and-observability-8417d1952e1c

Wrapping Up

In Summary

- Being fully cloud native requires new tech and new workflows
  - Lots to be learned from full cycle development

- Creating a supporting cloud platform is essential
  - Container orchestration
  - Progressive delivery
  - Edge management
  - Observability

- Consciously design your platform & watch for antipatterns

thenewstack.io/learning-kubernetes-the-need-for-a-realistic-playground/

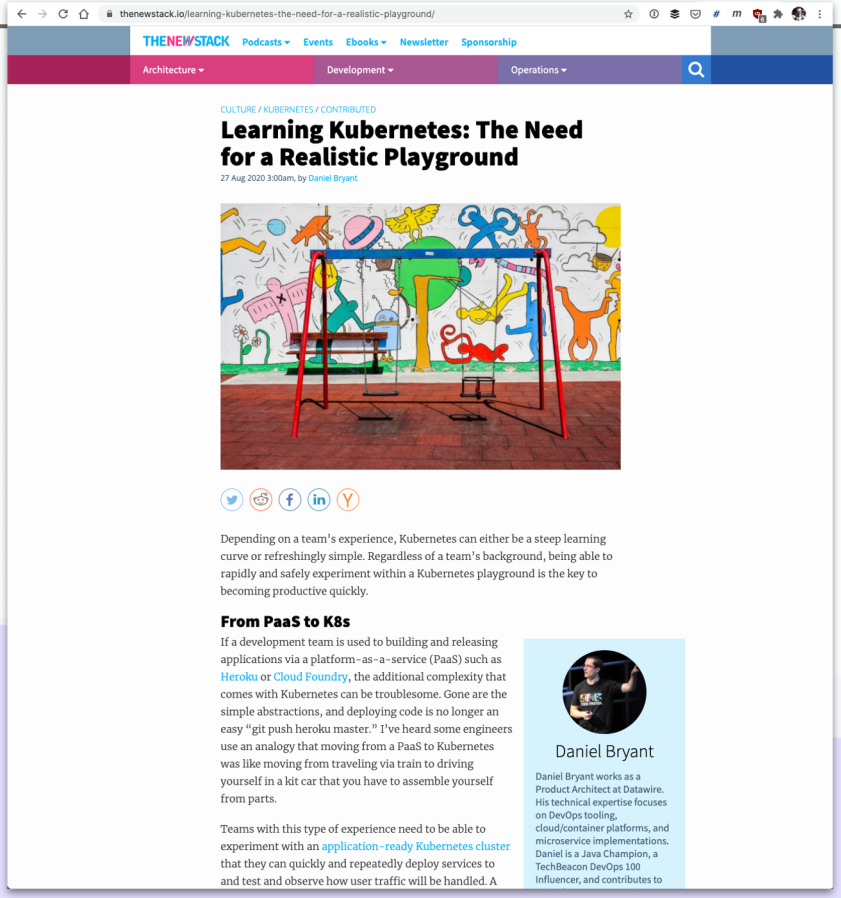

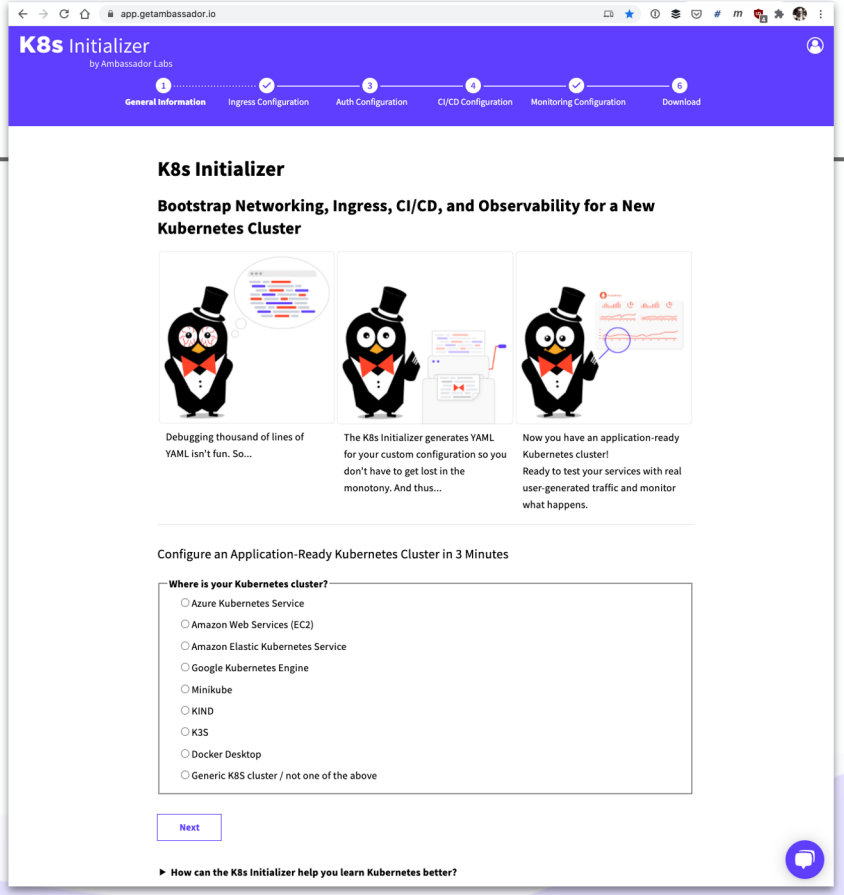

app.getambassador.io/

DATAWIRE

@danielbryantuk

Learning More...

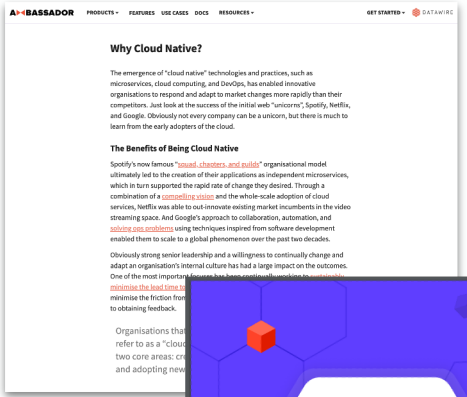

Read “Building a Kubernetes Platform”:
https://www.getambassador.io/learn/building-kubernetes-platform/

Subscribe to podcasts:
https://www.getambassador.io/podcasts/

Follow us on Twitter:
https://twitter.com/getambassadorio

Ambassador CNCF Incubations proposal:
https://github.com/cncf/toc/pull/435

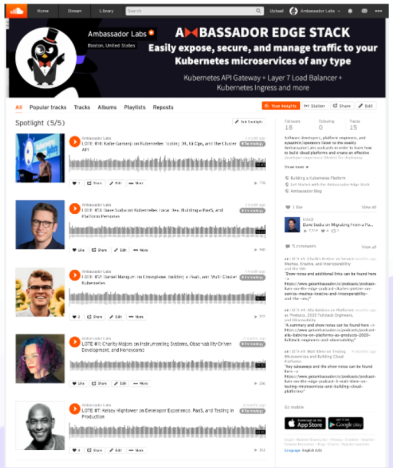

DATAWIRE
@danielbryantuk