# Airbnb's Istio Journey

Weibo He & Stephen Chan, 02/15/2021, IstioCon 2021

# The beginning

Where it all started

# 2013 - SmartStack Inception

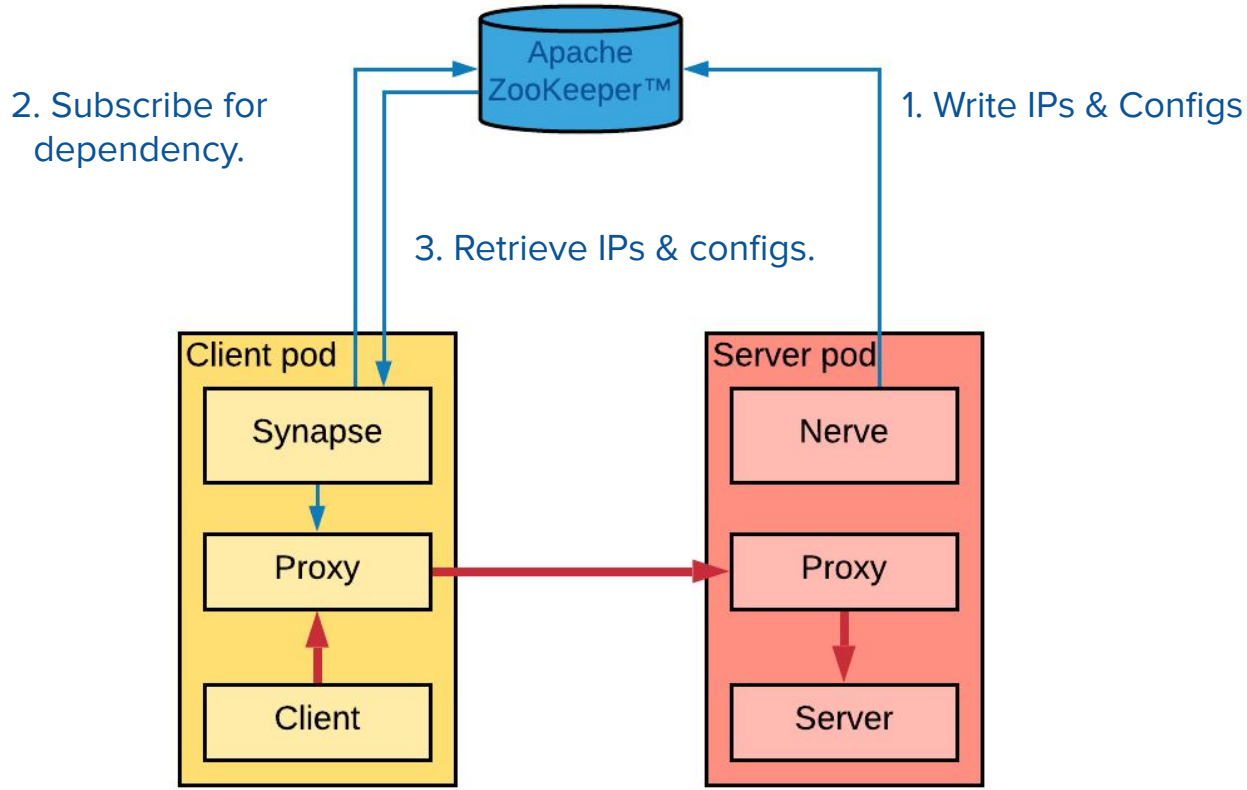

# 2018 - Scalability Issues

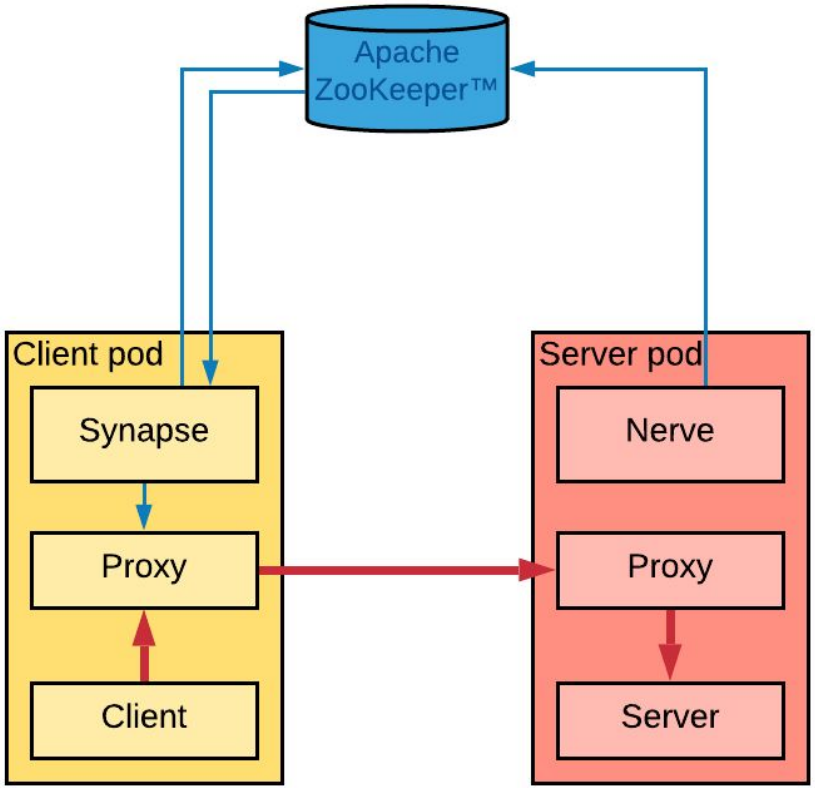

# 2019 - The search began

- Performance & Scalability
- Security - mTLS, easy cert rotation
- Data plane - preferably Envoy®
- Rich mesh features
- Works for K8s® & EC2

# Betting on Istio

## Why did we choose Istio?

# Early 2019 - Enter Istio

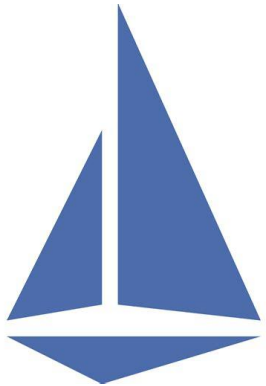

What we liked

- Security
- Envoy Data plane
- Traffic Management
- Resilience & Observability
- Active Community

Question Mark

- Performance & Scalability
- Multi Cluster Support
- VM Support

# Control Plane Performance

## Alleviating Concerns

- Examined control plane perf under varying # of proxies, namespaces, degree of connectivity & rate of changes
- Validate that the number of managed k8s clusters is not a control plane scaling factor for Istio 1.3

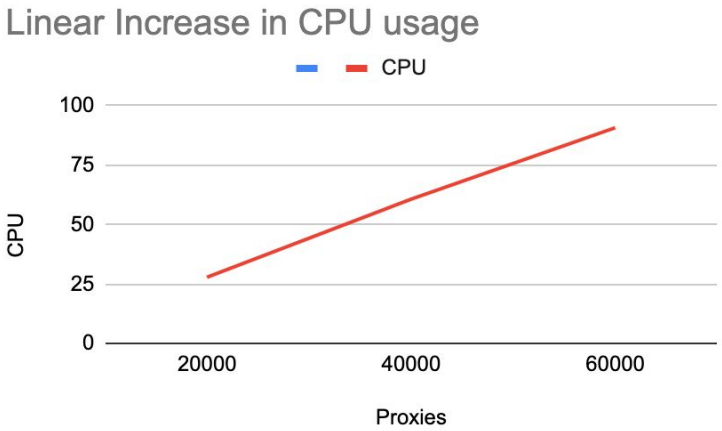


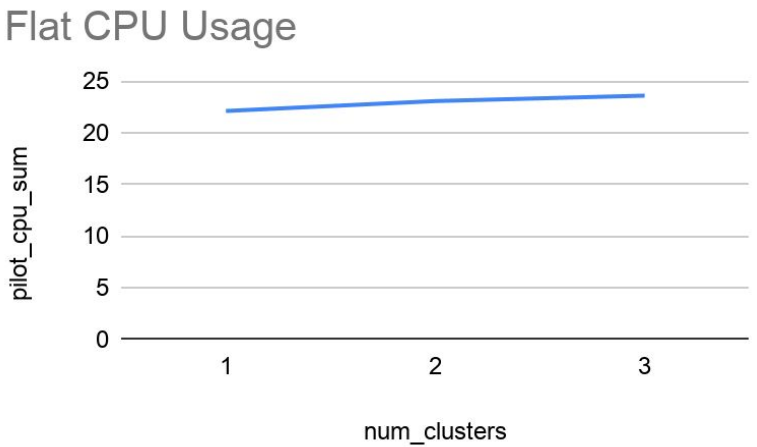

# Multi Cluster Support

## Alleviating Concerns

- We like external Istiod deployment model:
  - Tight access control for control plane.
  - Isolation from data plane workloads.
  - Ease of operation.
- Problems we ran into on Istio 1.5
  - Multi cluster DNS
  - Multi cluster CA
  - Multi cluster sidecar injection

# EC2 Support

## Alleviating Concerns

- VM support was primitive.
- We evaluated a few ways to support EC2.
  - Zookeeper plugin for Istio
  - Custom sidecar + controller to update ServiceEntry
- Discontinued in favor of community solution

# Airbnb's Istio Setup

## Across multiple clusters & environments

# Architecture

## Single Region

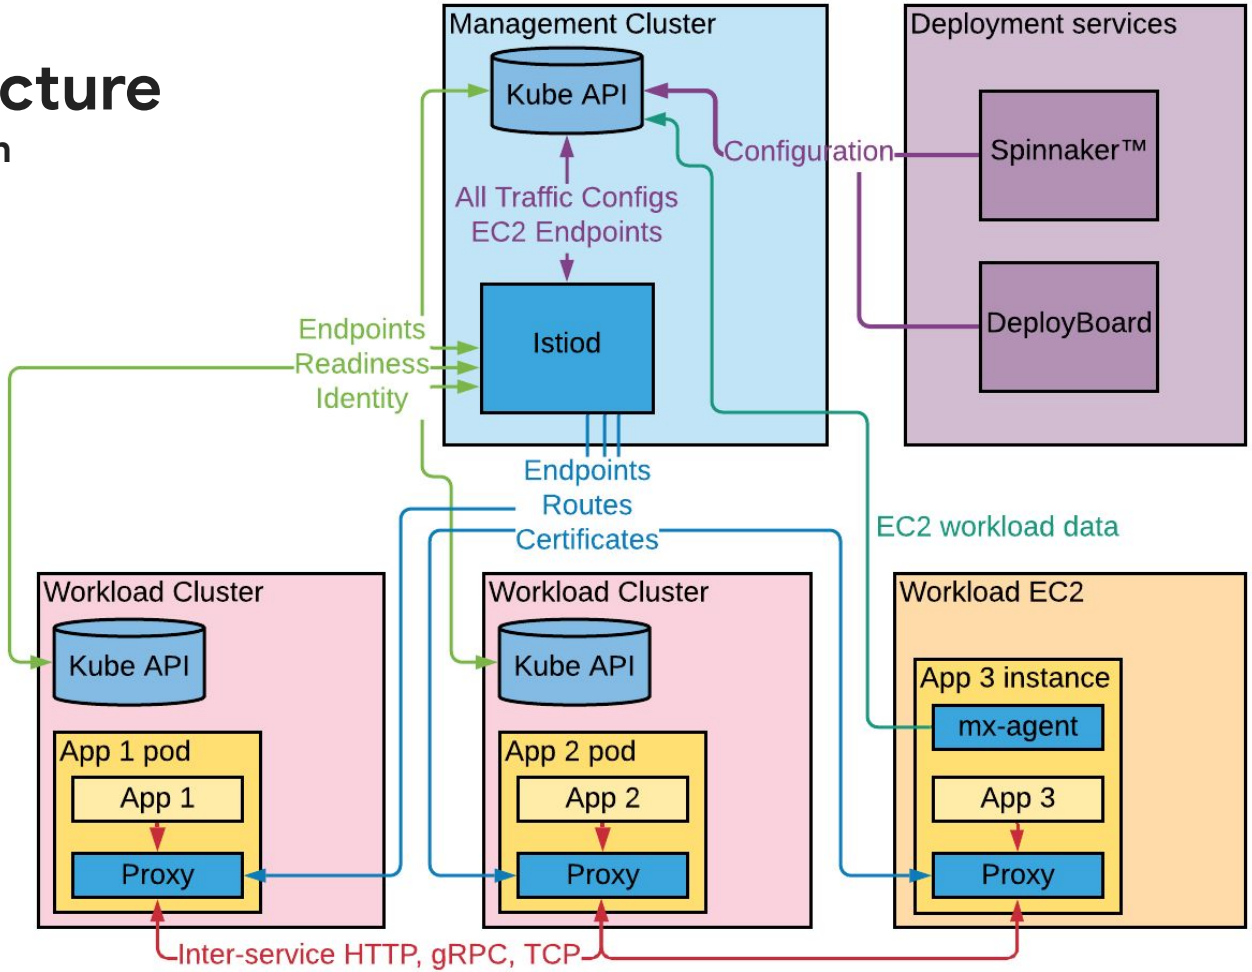

# Architecture

## Multi-region

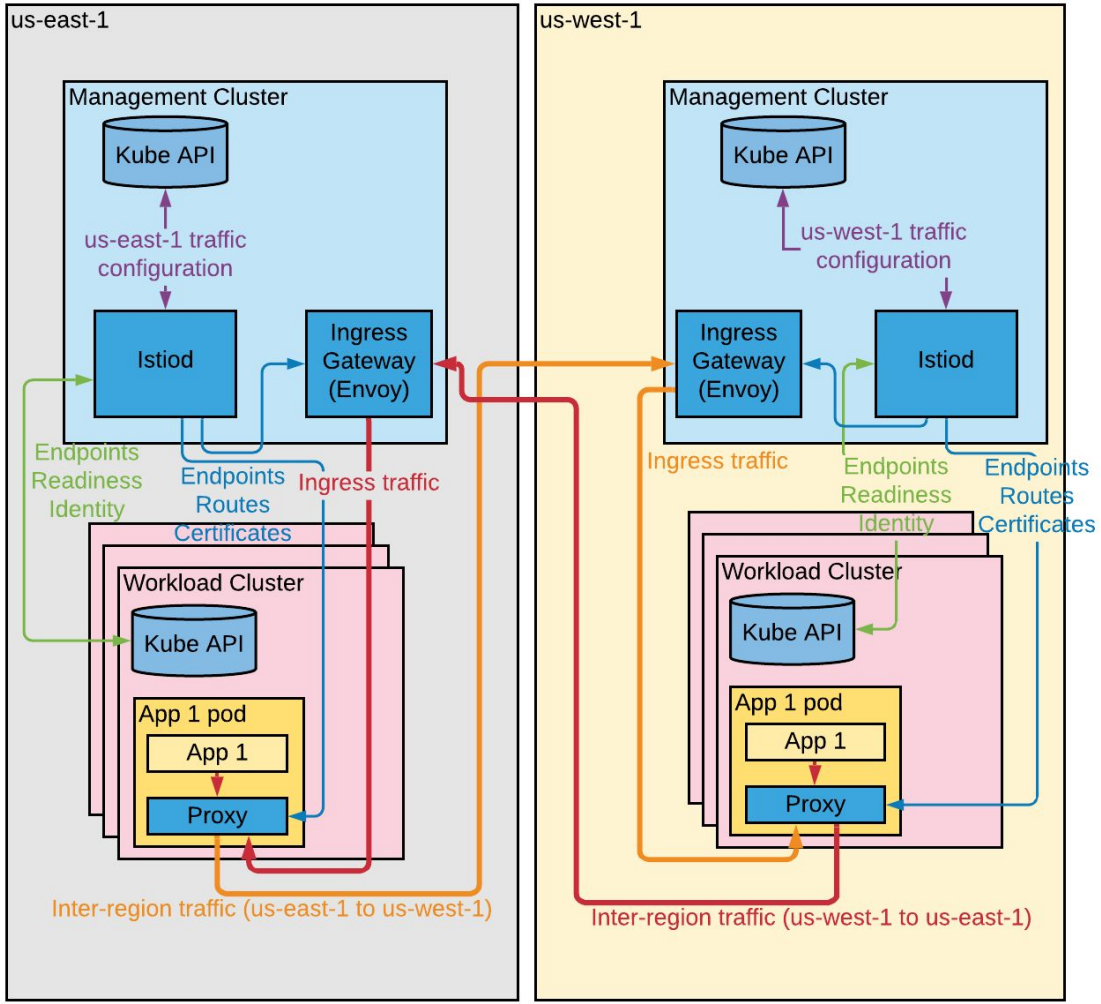

# CA Hierarchy

## One root certificate per trust domain

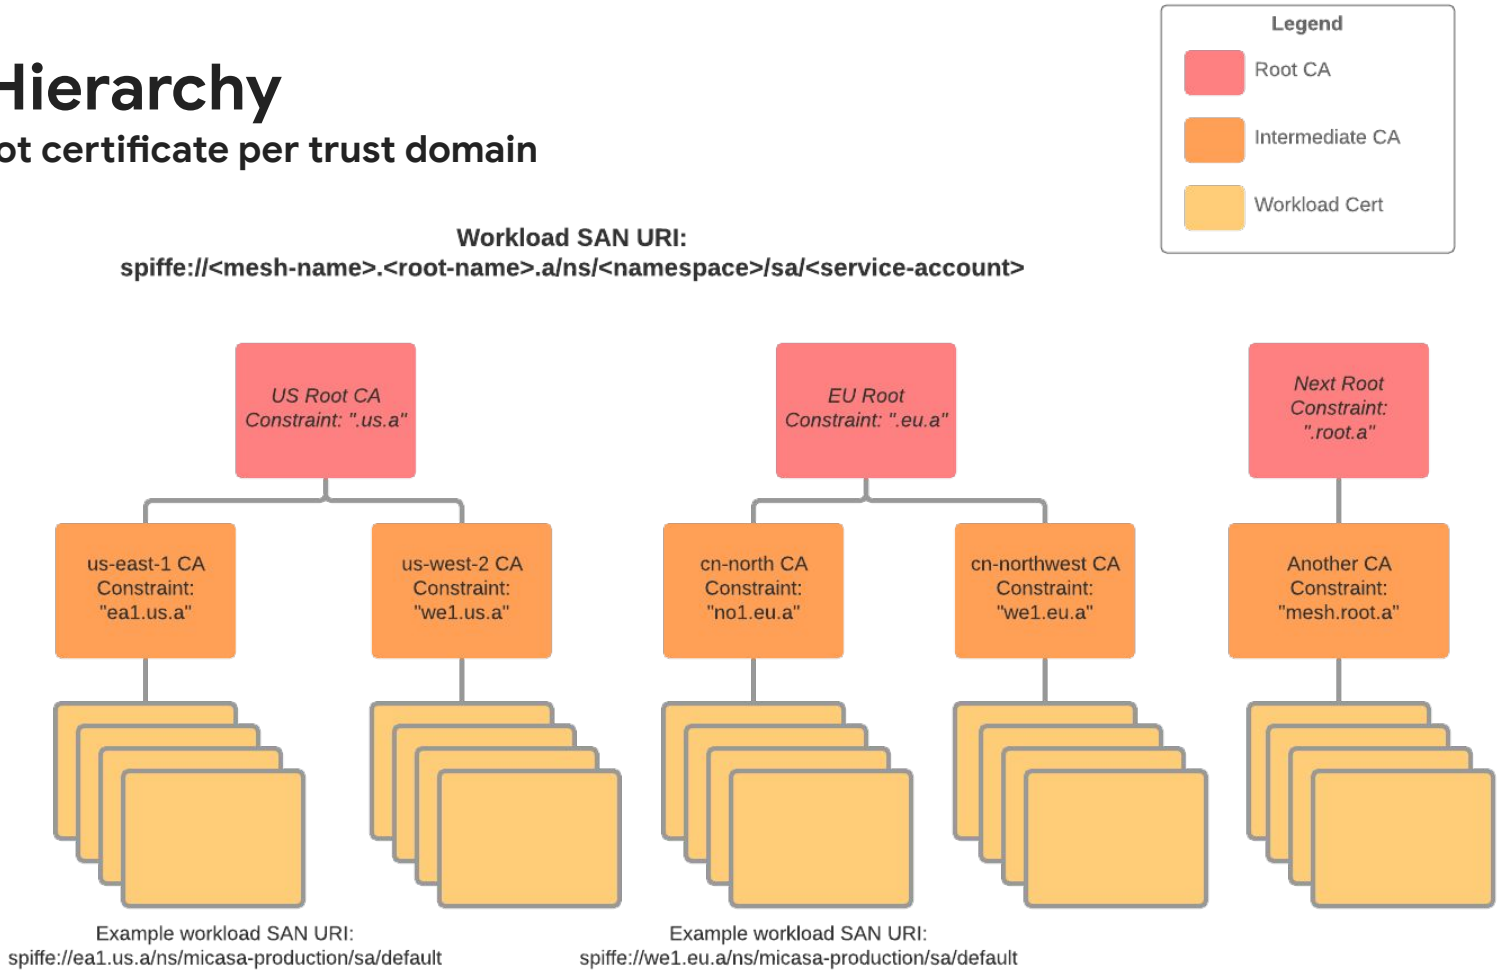

# The Great Migration

Changing inter-service communication on the fly

# Migration Requirements

**Migrations are hard**

Safe

- Edge by edge
- Percentage based gradual rollout
- Instant Rollback

Easy

- No code change & minimum config changes
- Automation whenever applicable

# URL controlled traffic shifting

## Smartstack & Istio side by side

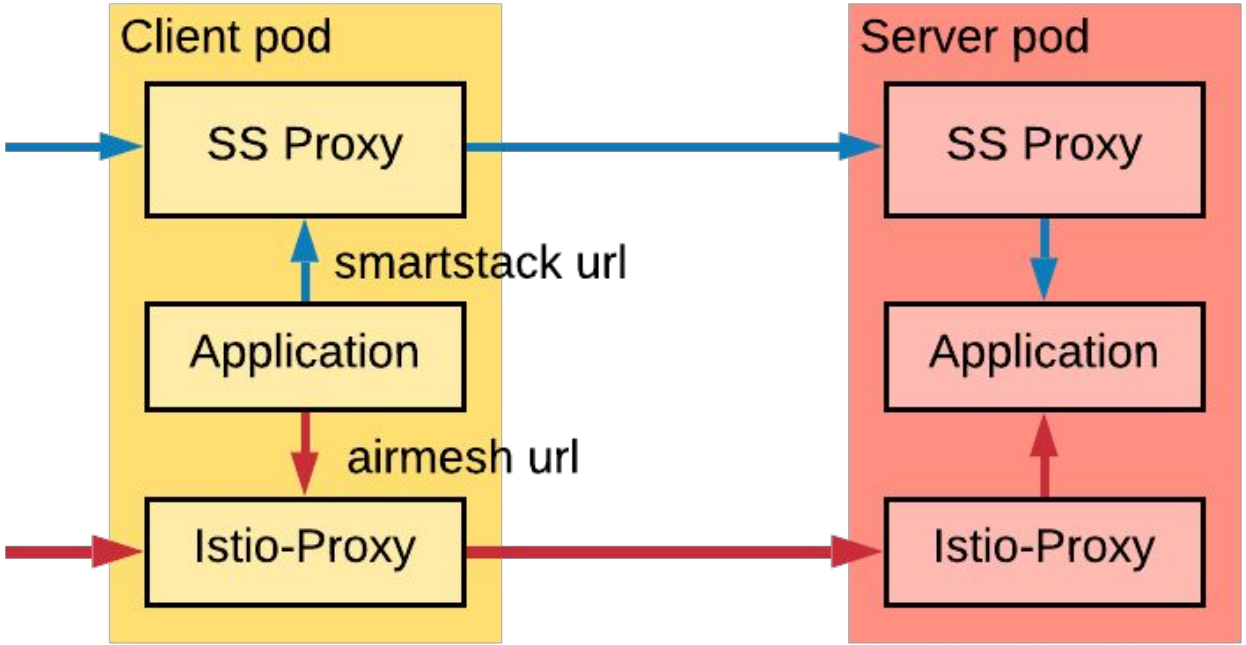

# Upgrading Istio

Keeping up with quarterly release

# Overview

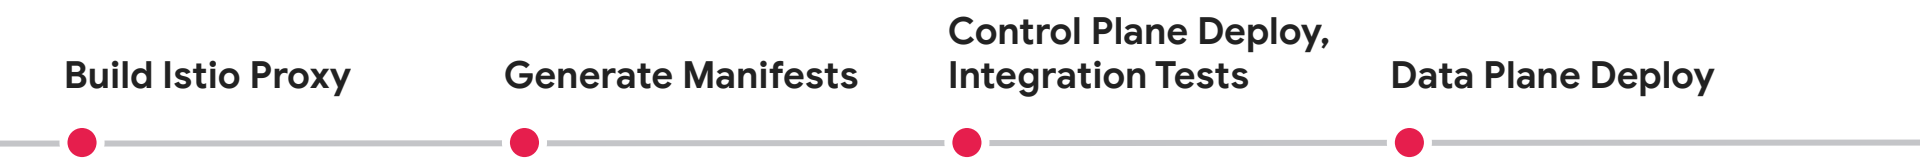
Build Istio Proxy
Generate Manifests
Control Plane Deploy,
Integration Tests
Data Plane Deploy

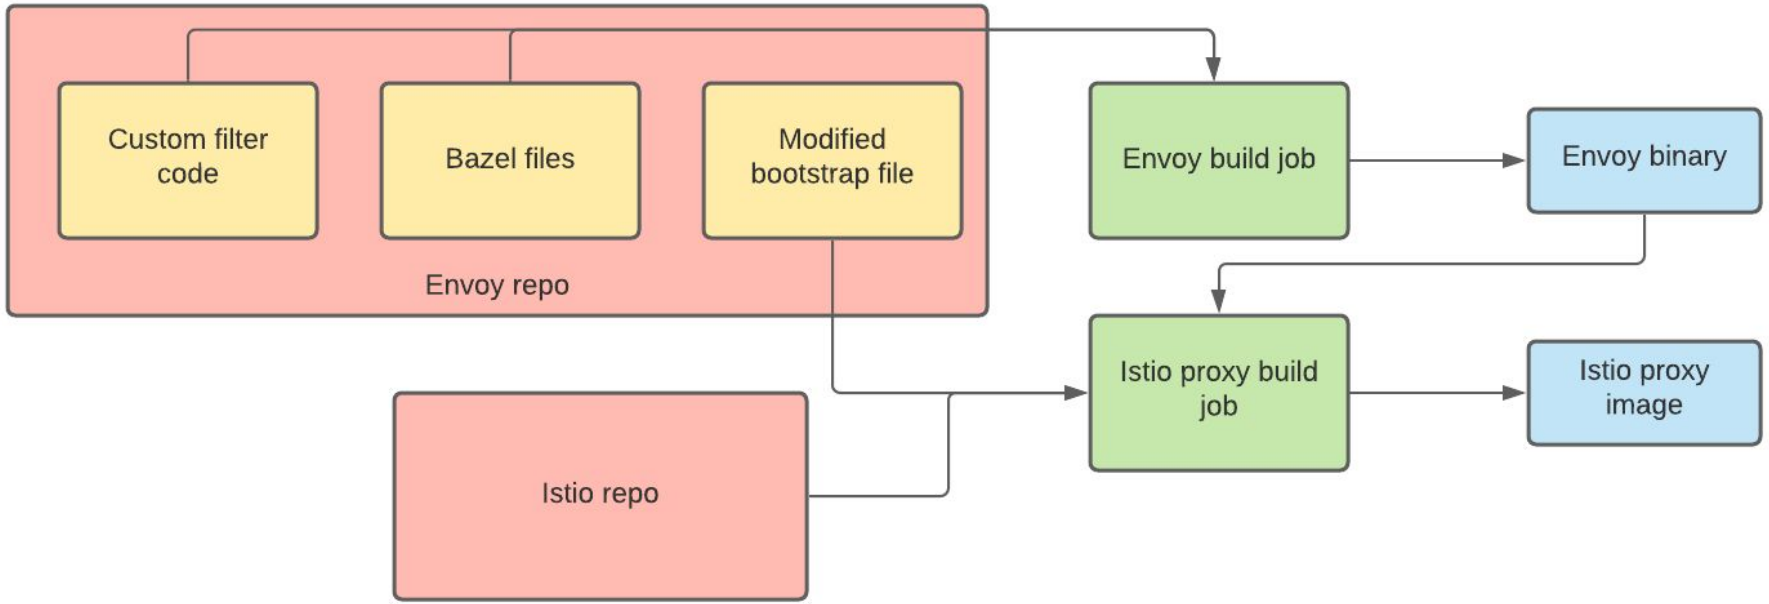


**Build Istio Proxy**

Generate Manifests

Control Plane Deploy, Integration Tests

Data Plane Deploy

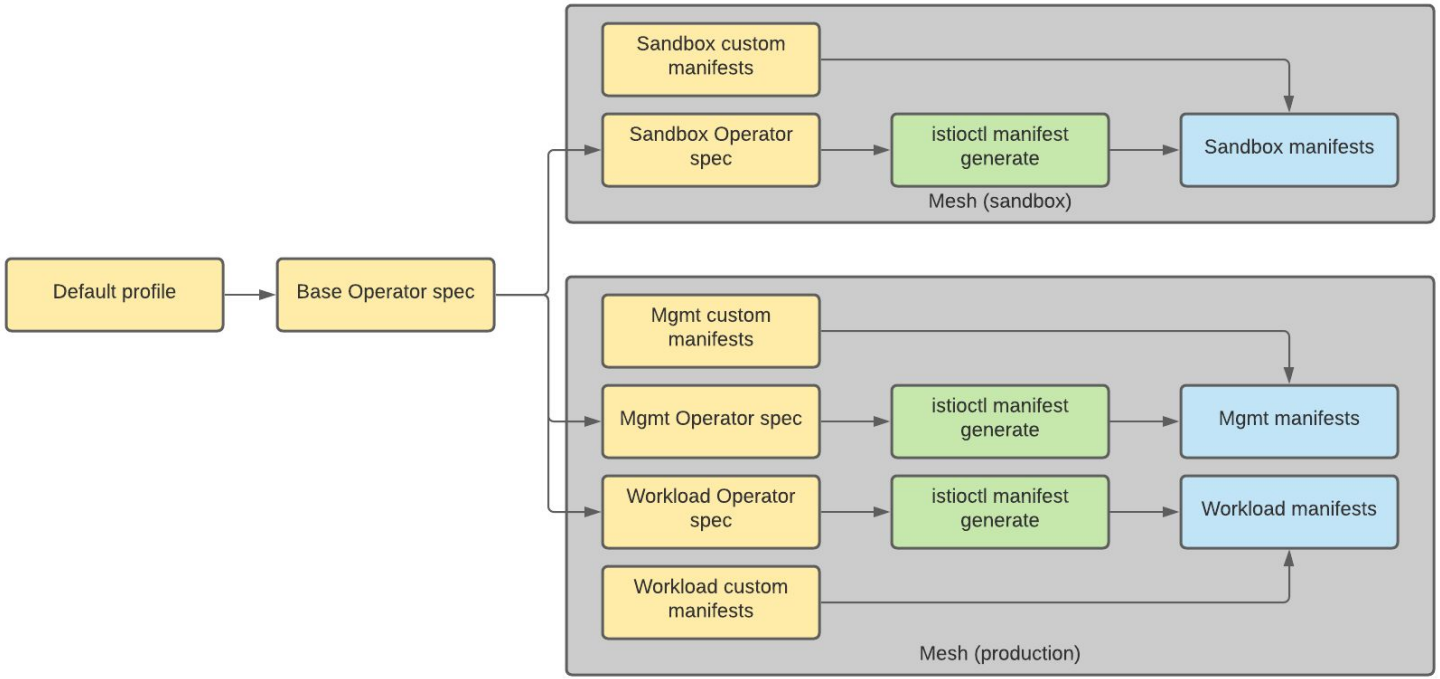


Build Istio Proxy

## Generate Manifests

Control Plane Deploy, Integration Tests

Data Plane Deploy

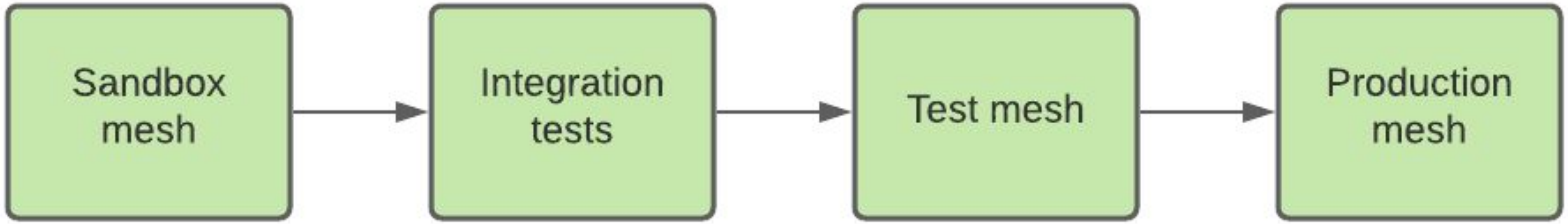


Build Istio Proxy — Generate Manifests — **Control Plane Deploy, Integration Tests** — Data Plane Deploy

# Testing: Cases & Setup

- Data Plane resource consumption & latency overhead
- AuthorizationPolicy
- Resilience Features (retries, timeout, circuit breaker, etc)
- Load Balance, locality aware routing
- Traffic Management features
- Fortio for load testing

# Testing: Issues Detected

- Incorrectly skipped XDS push when Istio CRDs are changed
- CDS is not updated as destination changes in VirtualService
- Changes in delegate virtual service do not take effect

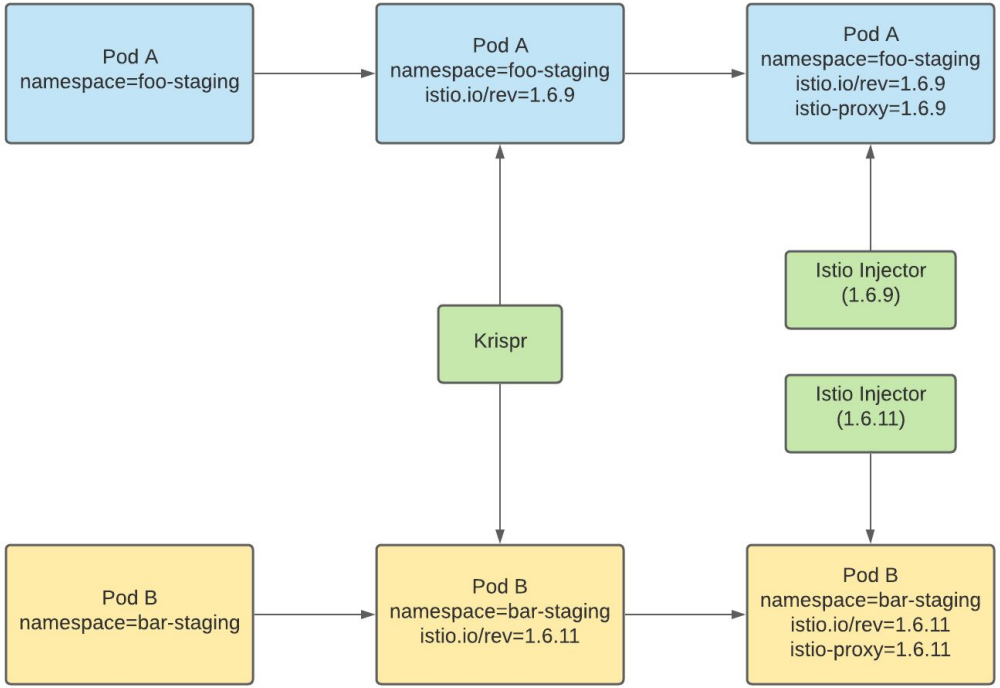


Build Istio Proxy · Generate Manifests · Control Plane Deploy · **Data Plane Deploy**

Krispr blog post: https://medium.com/airbnb-engineering/a-krispr-approach-to-kubernetes-infrastructure-a0741cff4e0c

# Mesh Expansion

Supporting Istio on VMs

# Mesh Expansion: Istio 1.5

**Alpha**

| **Feature** | **Istio on k8s** | **Istio on VMs** |
|---|---|---|
| Installing proxy | Sidecar injection | Download .deb |
| Configuring proxy | Injection template, Proxyconfig (defaults), Annotations (customization) | Edit sidecar.env |
| PKI bootstrap | Namespace controller (CA), service account token (attestation) | transfer CA, token from admin env |
| Health checking app | Container probes | **???** |
| Register service instance | Endpoints controller updates automatically | **Create or edit ServiceEntry (per instance)** |

# Mesh Expansion: Istio 1.8

**Almost Beta!**

| Feature | Istio on k8s | Istio on VMs |
|---|---|---|
| Installing proxy | Sidecar injection | Download .deb or .rpm |
| Configuring proxy | Injection template, Proxyconfig (defaults), Annotations (customization) | Edit sidecar.env or istioctl x workload entry configure |
| PKI bootstrap | Namespace controller (CA), service account token (attestation) | transfer CA, token from admin env |
| Health checking app | Container probes | WorkloadGroup Readiness probes |
| Register service instance | Endpoints controller updates automatically | WorkloadEntry, auto registration with WorkloadGroup |

# Mesh Expansion @ Airbnb

## Customizations

| Feature | Istio on k8s | Istio on VMs |
|---|---|---|
| Installing proxy | Sidecar injection | ~~Download .deb or .rpm~~<br>Custom artifact |
| Configuring proxy | Injection template, Proxyconfig (defaults), Annotations (customization) | ~~Edit sidecar.env or istioctl x workload entry configure~~<br>Template proxy run script |
| PKI bootstrap | Namespace controller (CA), service account token (attestation) | ~~transfer CA, token from admin env~~<br>VM requests CA, token from k8s (future: IdentityProvider?) |
| Health checking app | Container probes | WorkloadGroup Readiness probes |
| Register service instance | Endpoints controller updates automatically | WorkloadGroup, auto registration |

# Mesh Expansion @ Airbnb

## Additional Priorities

- Data plane gradual rollout
- Abstracting WorkloadGroup details
- Safe edge-by-edge migration

# Future Usage & Takeaways

# Future Istio Usage @ Airbnb

- Advanced resiliency features
- Gateways for cross-mesh traffic
- TCP support
- gRPC support

# Takeaways

1. Istio's extensibility, broad feature support and scalability make it a great choice for Airbnb.
2. Consider using a management Istio cluster for a multi-cluster mesh.
3. Check manifest changes before upgrading the control plane.
4. Use canary control plane deployments and upgrade the data plane gradually.
5. Test features your services depend on before upgrading the data plane.
6. Migrate edge by edge using gradual traffic shifting.
7. Consider using auto registration and health probes for Mesh Expansion.
8. Engage with the community.

# We’re hiring!

**Open positions: airbnb.com/careers**