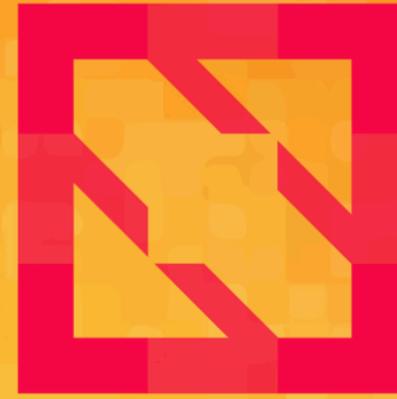




KubeCon



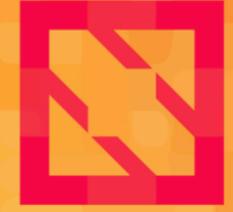
CloudNativeCon

North America 2019





KubeCon



CloudNativeCon

North America 2019

Building a Medical AI With Kubernetes and Kubeflow

Jeremie Vallee, Cloud Infrastructure Engineer, Babylon Health





KubeCon



CloudNativeCon

North America 2019

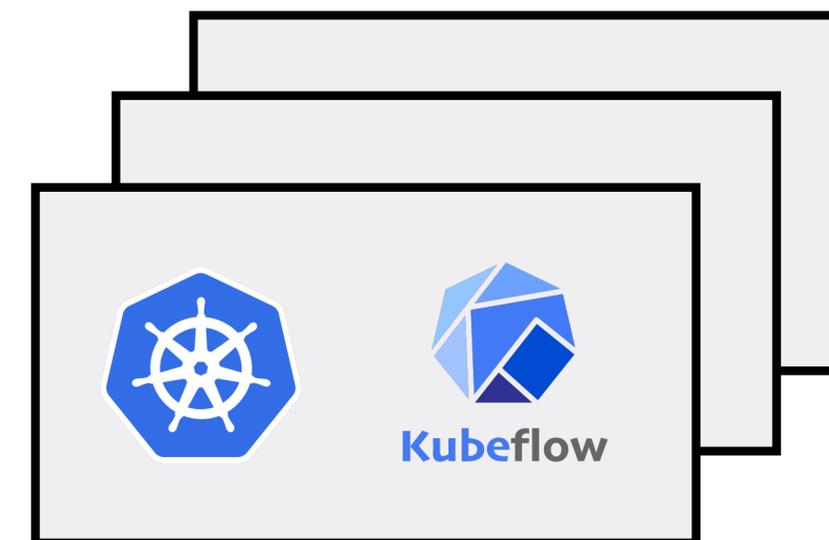
From



“Hi! 🙋”

Is anyone running a training job on the server? My notebook is super slow...”

To



Self-service
Multi Region
Multi Cloud

Research and Training Platform



KubeCon



CloudNativeCon

North America 2019



babylon

We believe it is possible to put an **accessible** and **affordable** health service in the hands of every person on earth.





KubeCon



CloudNativeCon

North America 2019



Accessibility

50% of the world population lacks access to essential health services, however **67% have access to mobile phones.**

In 2014:

Digital consultation with a doctor from your mobile phone.

Now:

We've had 2+ million digital consultations in the world.

Affordability?



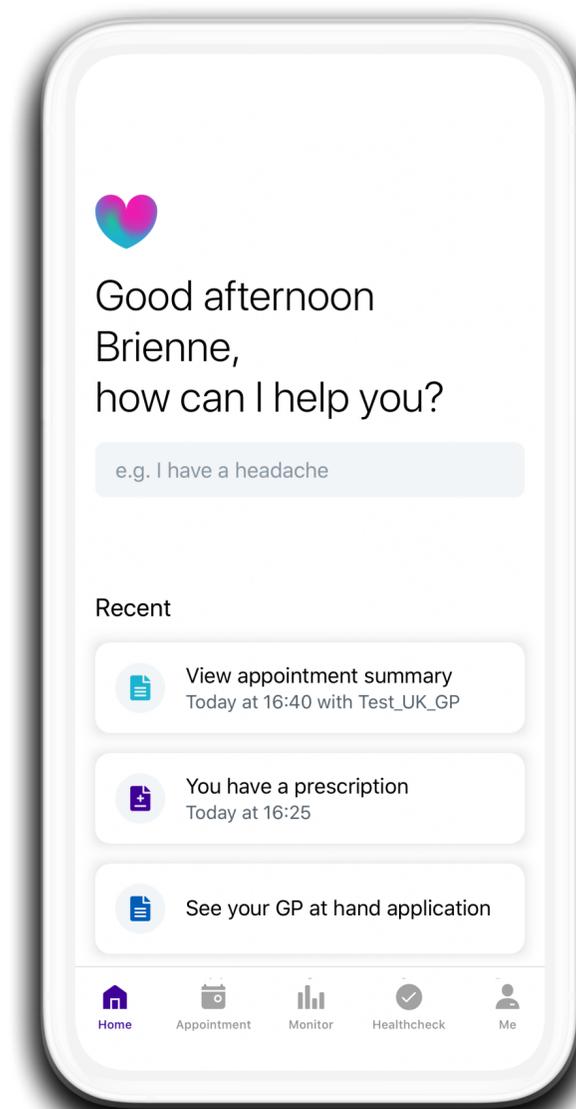
KubeCon



CloudNativeCon

North America 2019

Affordability



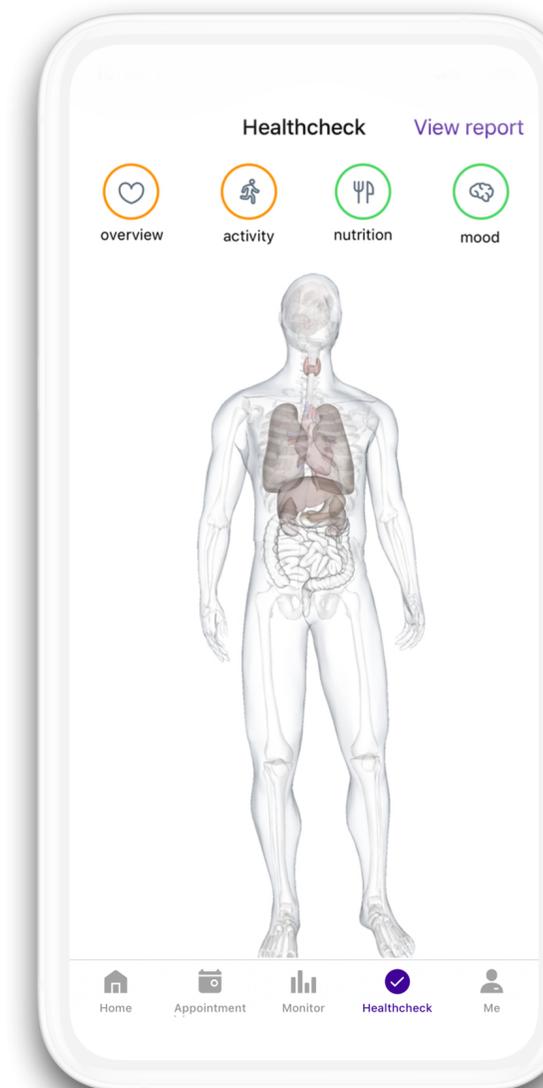
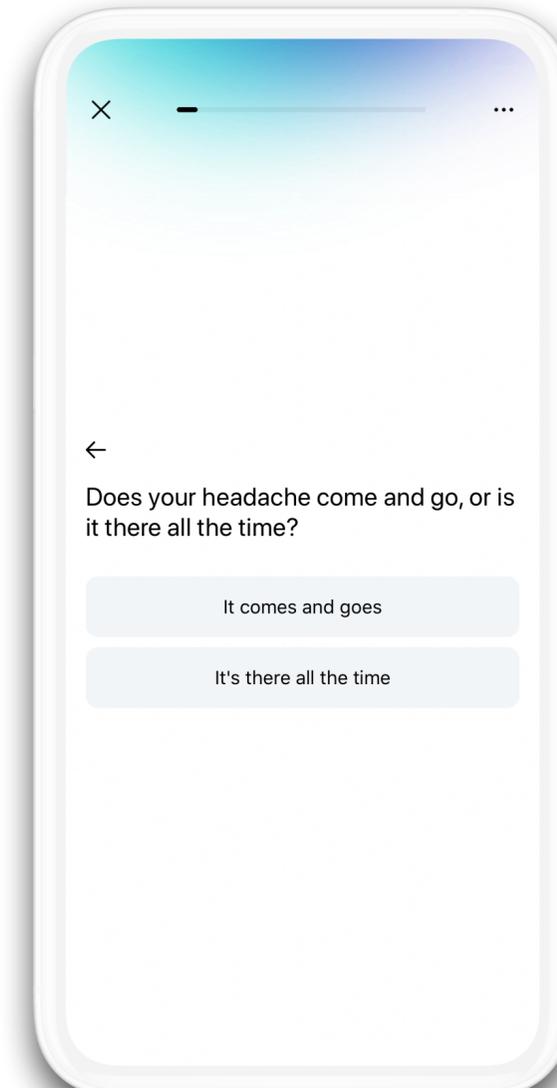
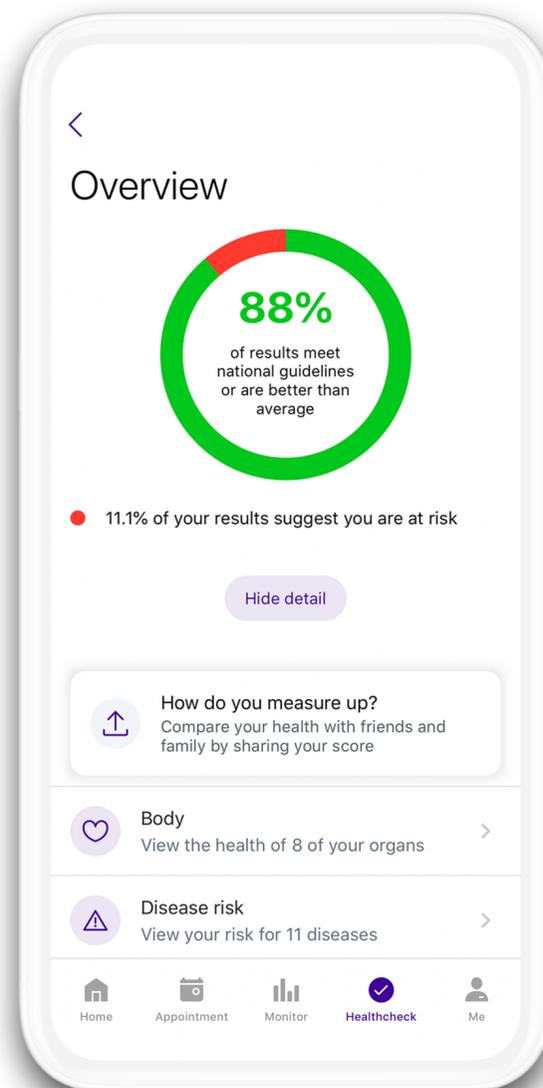
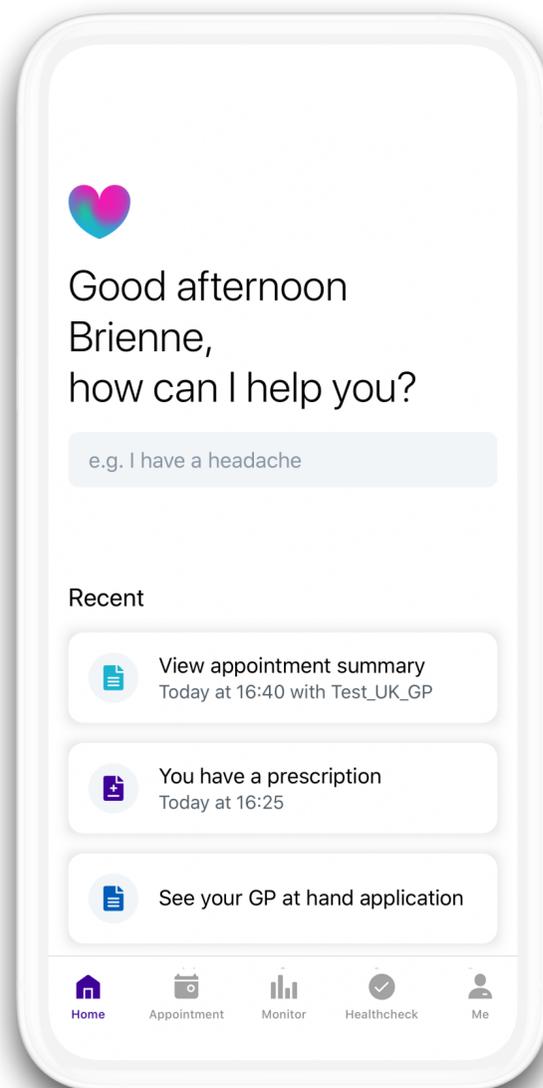


KubeCon

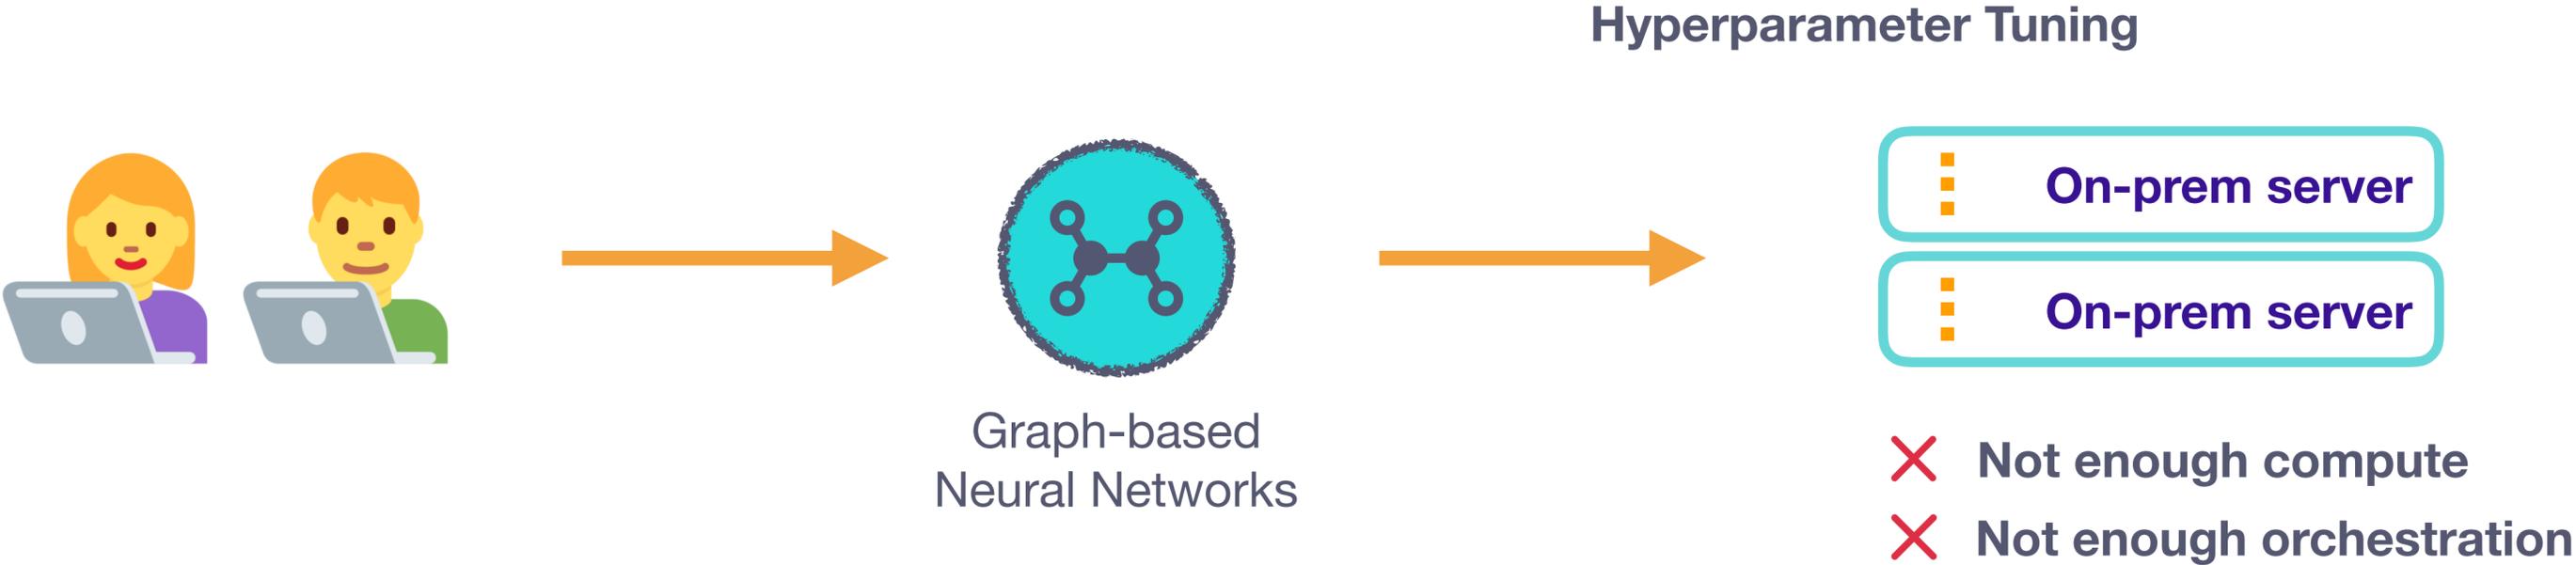


CloudNativeCon

North America 2019



How it started



Compute



HP Tuning
Orchestration

- On-prem server
- On-prem server

- Not enough compute
- Not enough orchestration

Hyperparameter Tuning with Katib

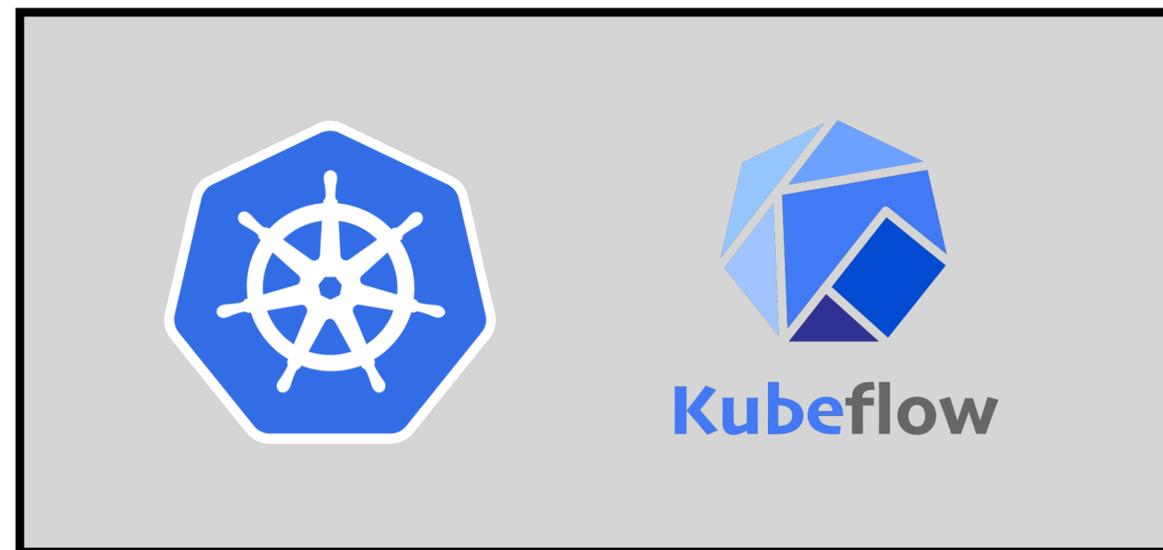


KubeCon



CloudNativeCon

North America 2019



1600 CPU
3.2TB RAM



Challenges & motivation



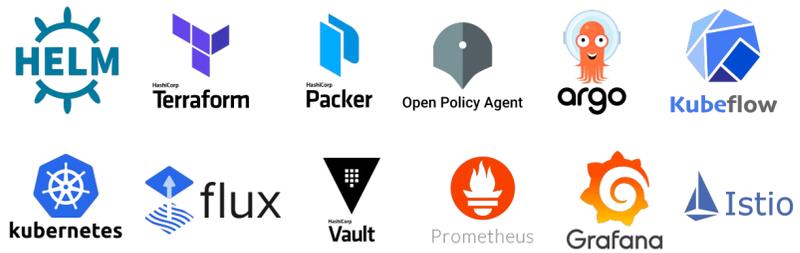
Growing company



Increasing need for | compute access
| AI Research & Engineering tooling



Global footprint



Let's talk about:

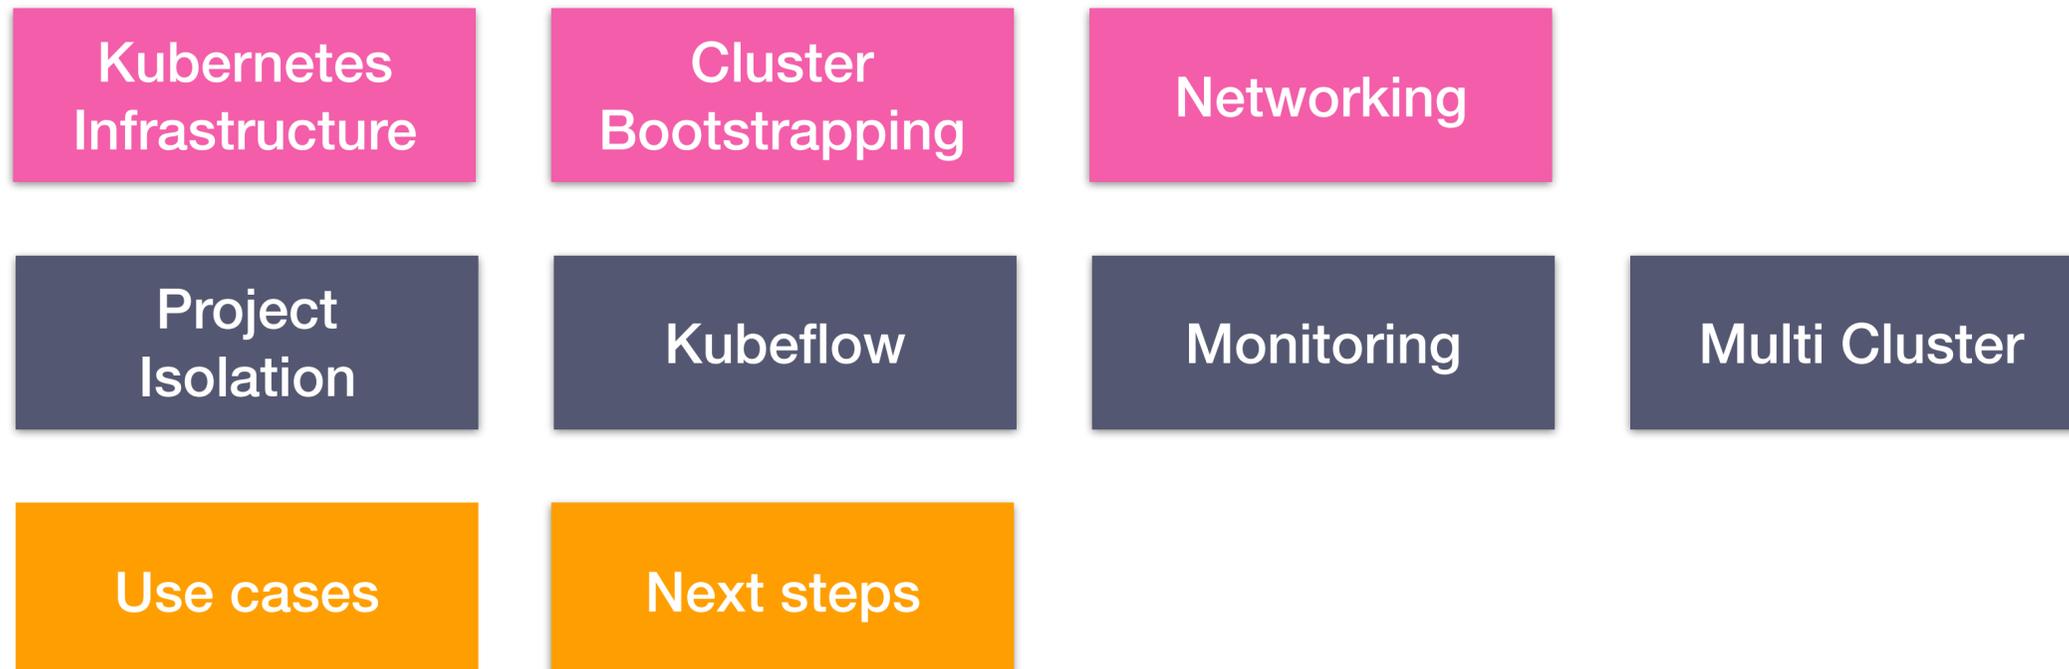


KubeCon



CloudNativeCon

North America 2019



Kubernetes Infrastructure



KubeCon



CloudNativeCon

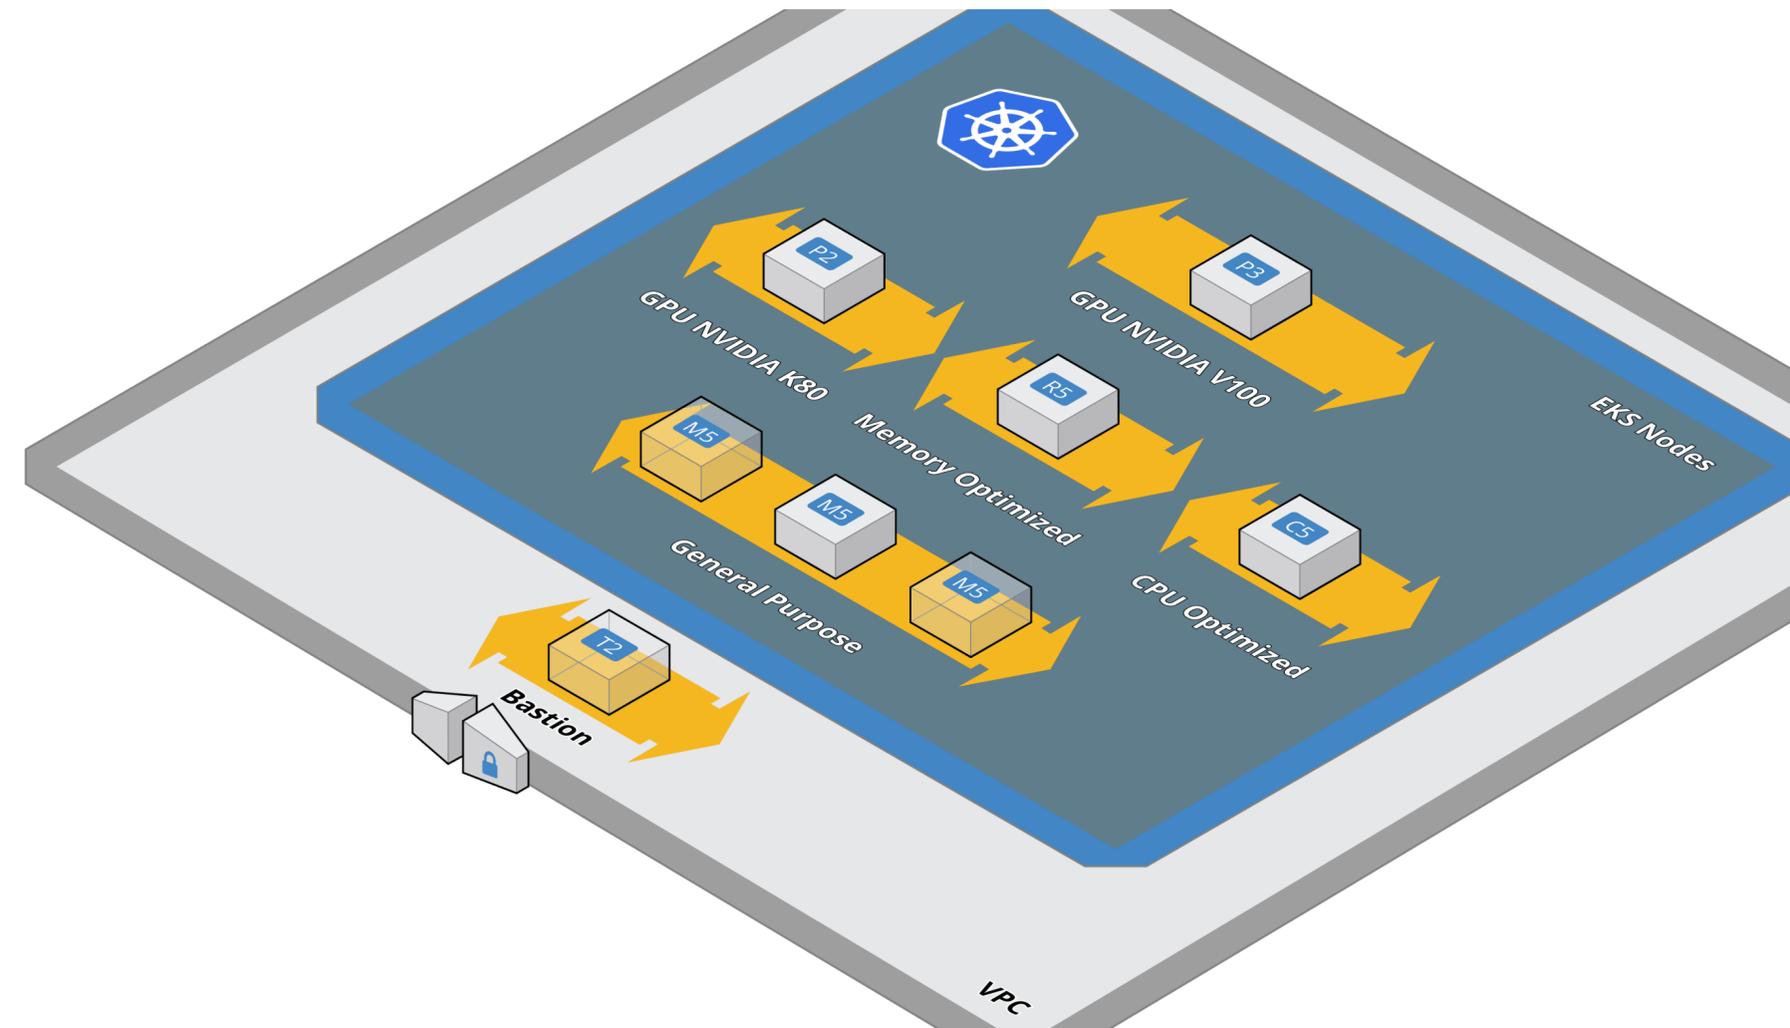
North America 2019

Prepare for different workload types

Harden your nodes

Make Kubernetes API private

Encrypt your nodes (root volumes + others)



GitOps Bootstrapping



KubeCon



CloudNativeCon

North America 2019



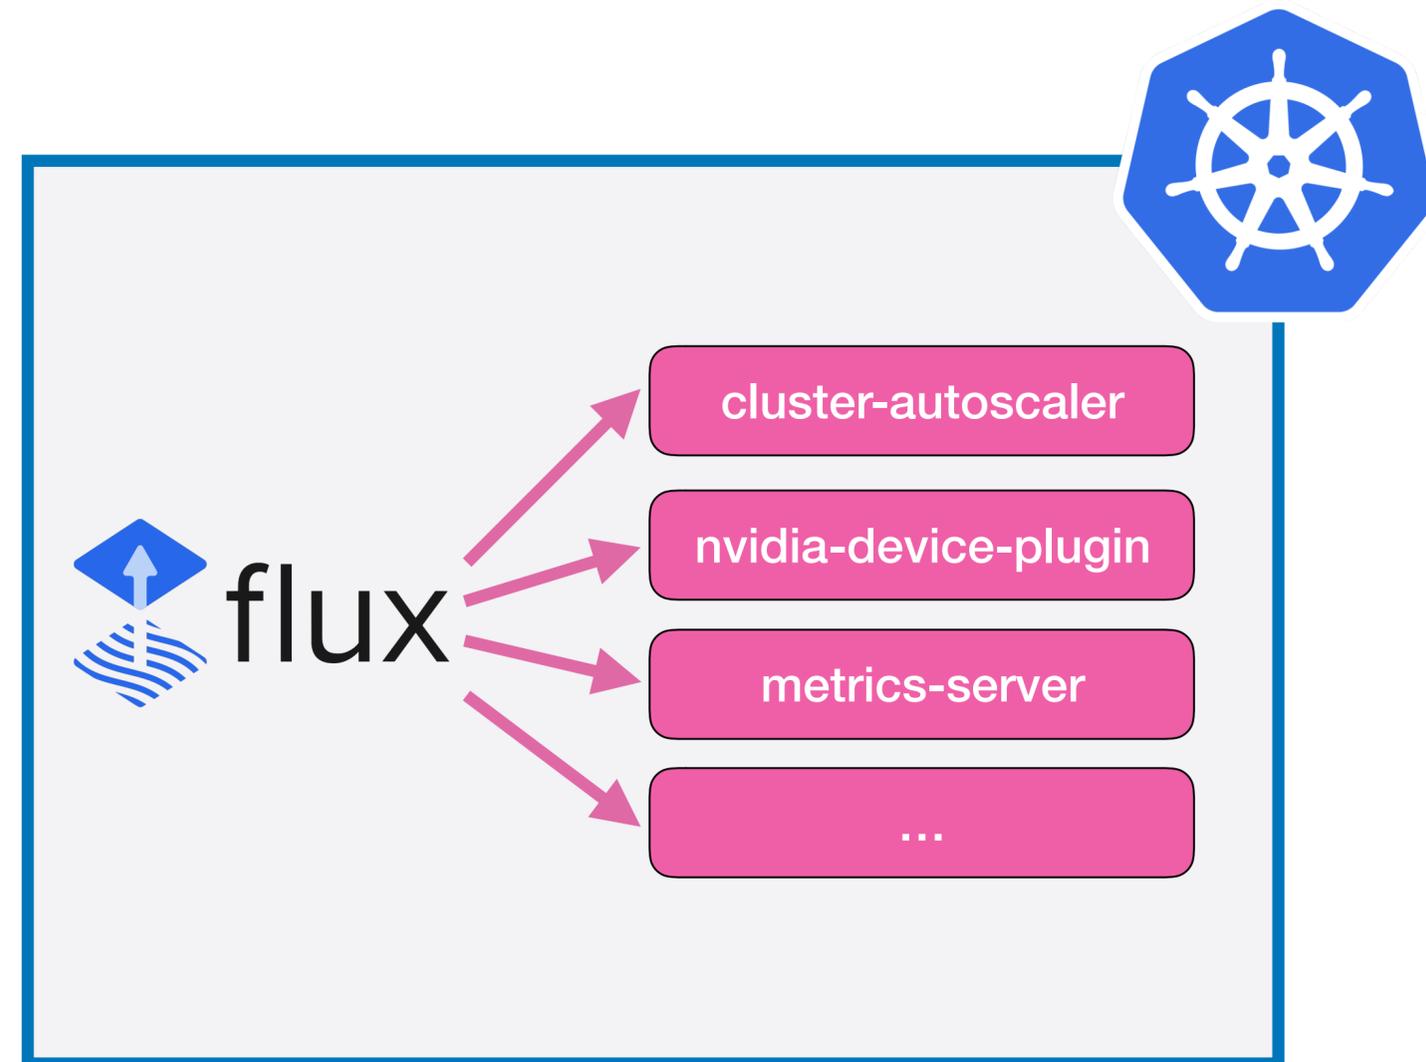
YAML



Source of Truth
Auditing



Pulls
for
change



Reduce human error 👍

Networking



KubeCon

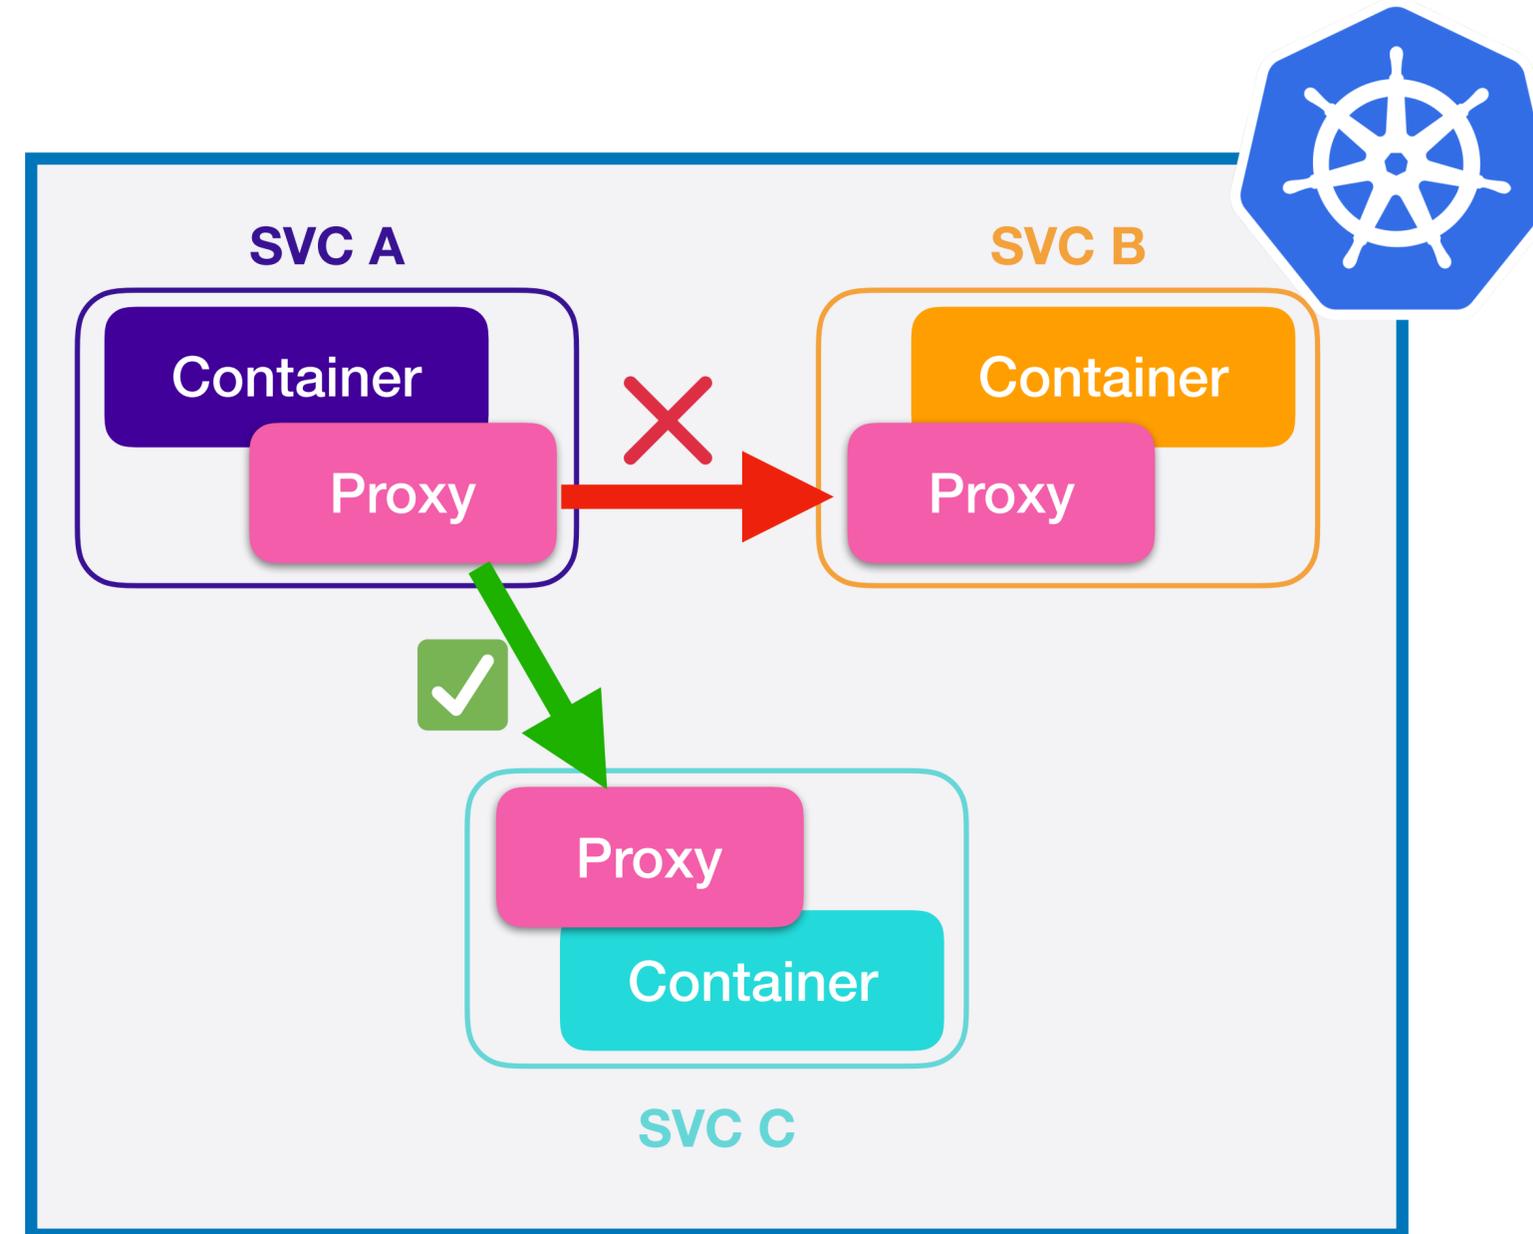


CloudNativeCon

North America 2019

```
apiVersion: rbac.istio.io/v1alpha1
kind: ServiceRole
metadata:
  name: c-access-role
  namespace: default
spec:
  rules:
    - methods:
      - '*'
      paths:
      - '*'
      services:
      - c.default.svc.cluster.local
```

```
apiVersion: rbac.istio.io/v1alpha1
kind: ServiceRoleBinding
metadata:
  name: bind-a-to-c
  namespace: default
spec:
  roleRef:
    kind: ServiceRole
    name: c-access-role
  subjects:
    - user: cluster.local/ns/default/sa/a
```



Zero-trust policy

Mutual TLS

JWT
check

Project Custom Resource



KubeCon



CloudNativeCon

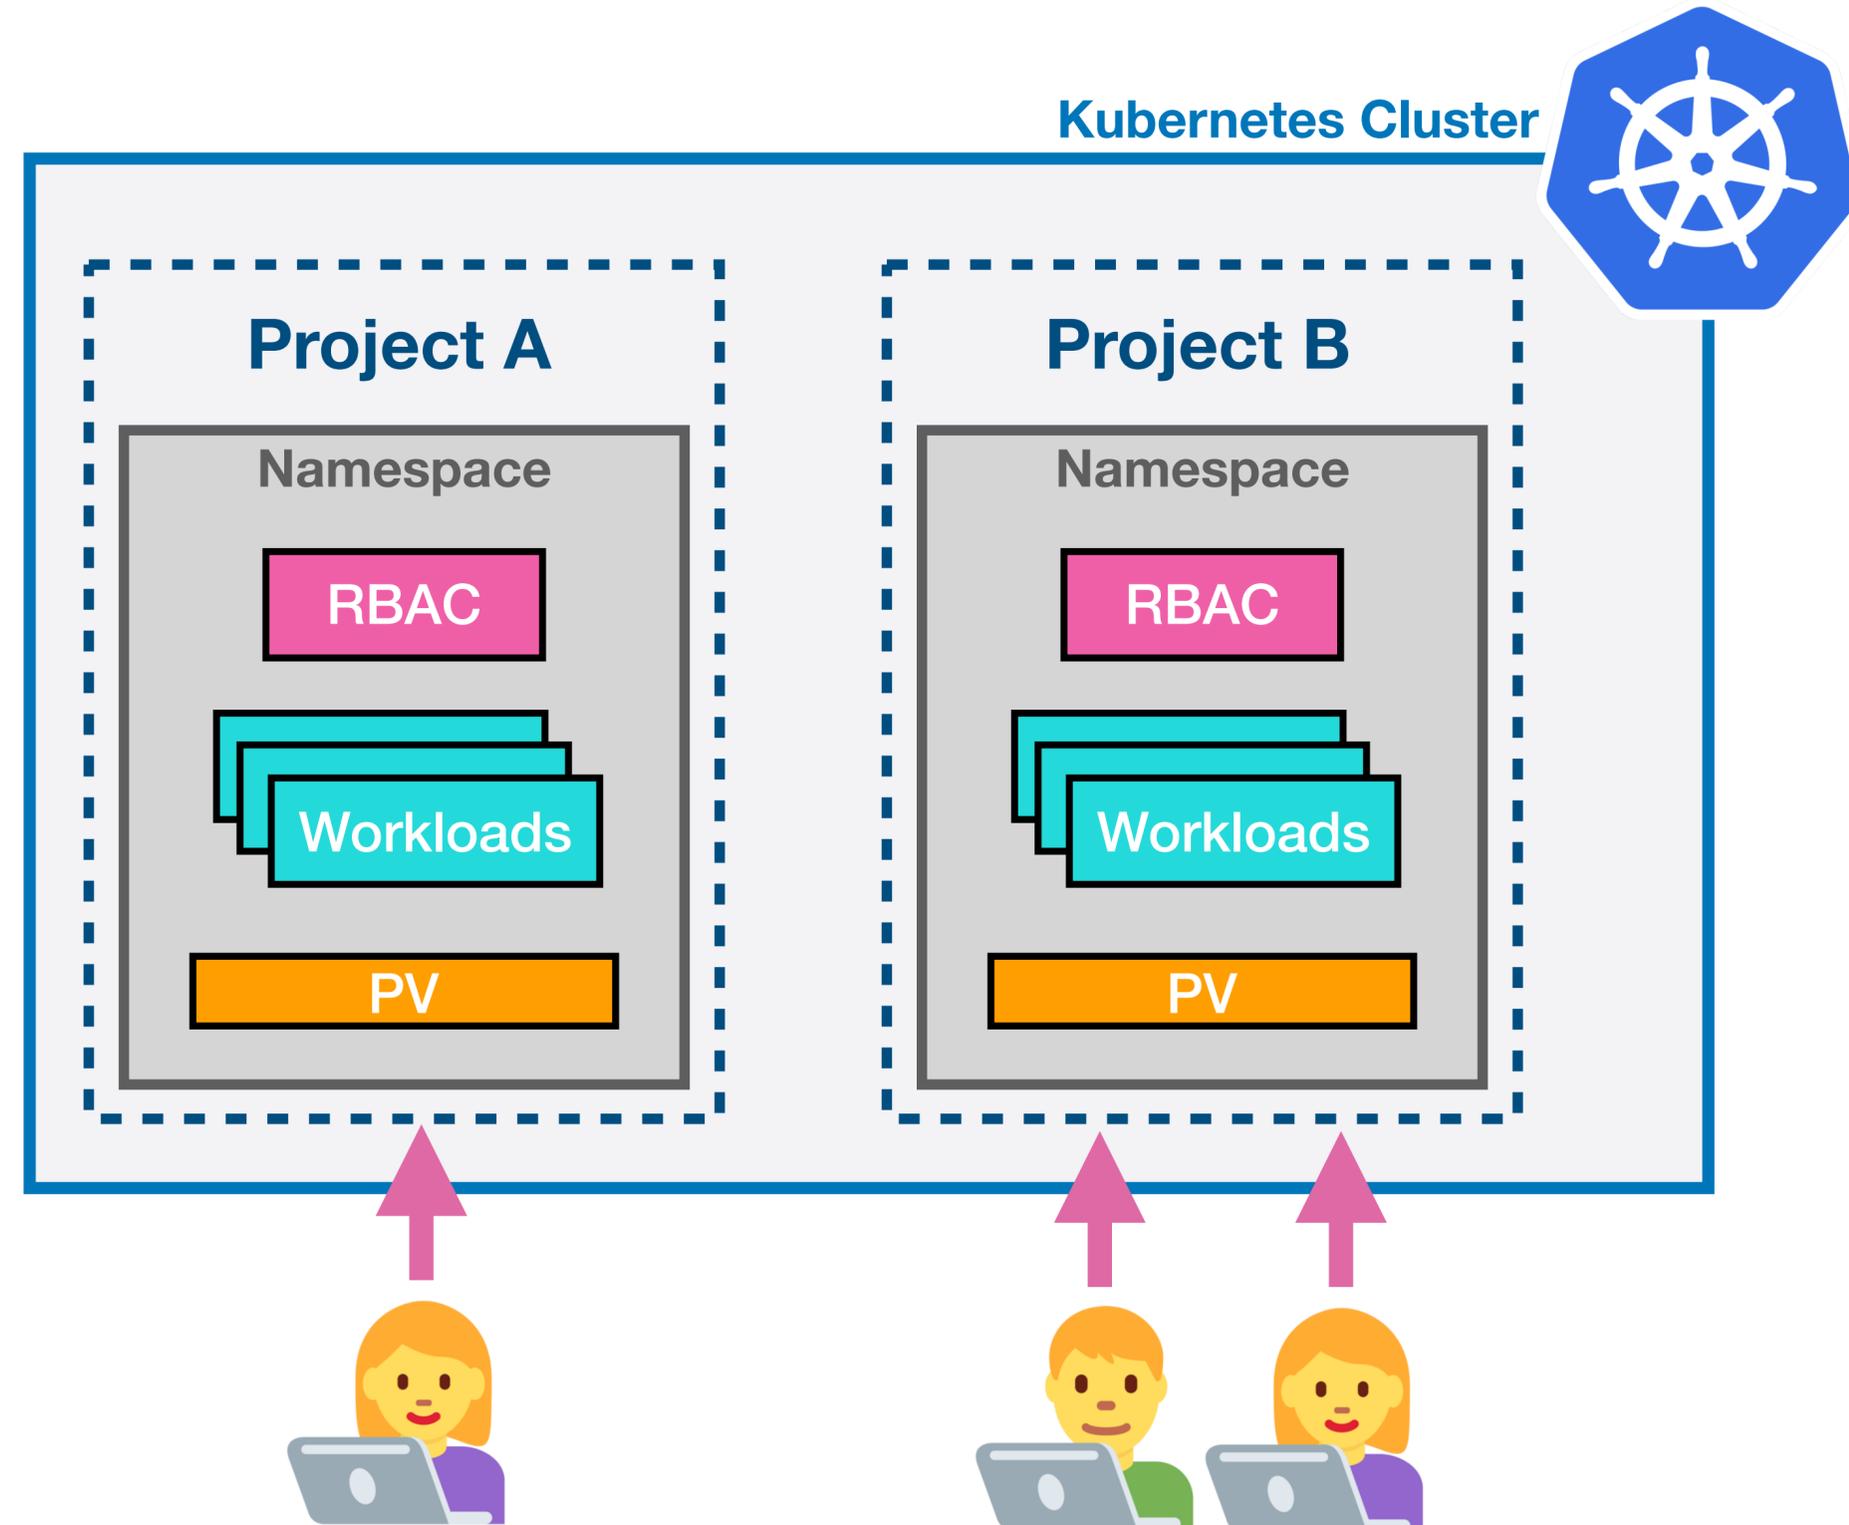
North America 2019

Goal:

- Isolate workloads
- Enable collaboration (multi-user)
- Provide shared volume per project
- Additional business logic and metadata

Solution:

- “*Project*” CRD and controller
- Integration with Kubeflow “*Profile*” CRD



Kubeflow

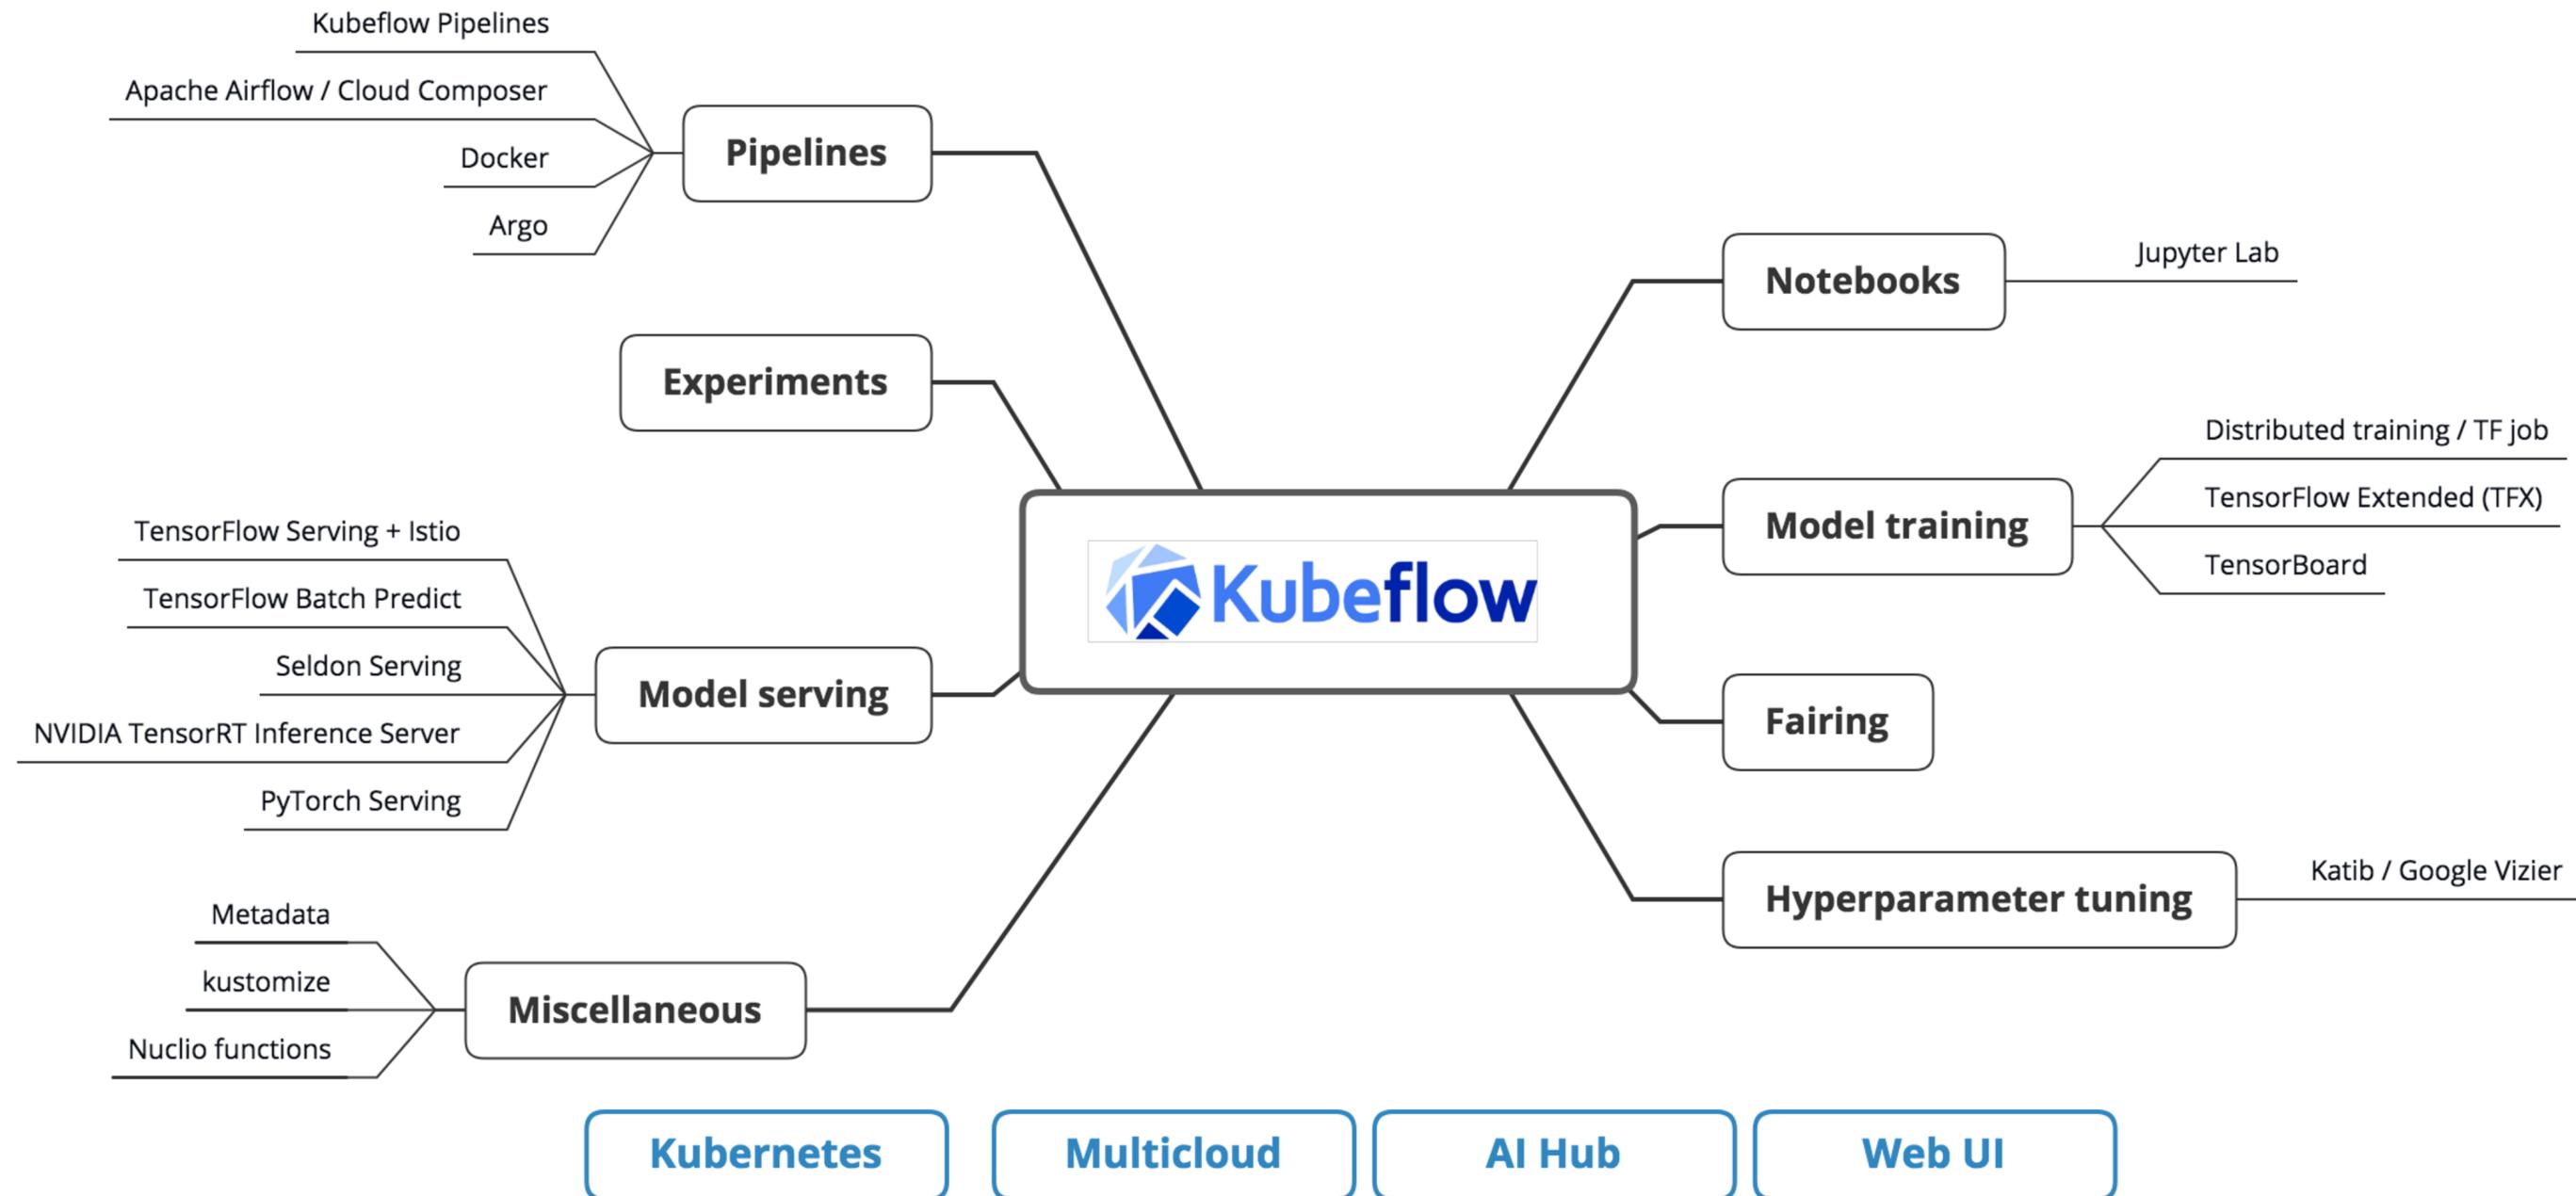


KubeCon



CloudNativeCon

North America 2019



Kubeflow

Modular: install what you need

Deployment: GitOps (Kustomize + Flux)



Katib

HP Tuning



**Notebook
Controller**

PYTORCH



Operators



argo

Workflows

Kubeflow + Istio + AuthN

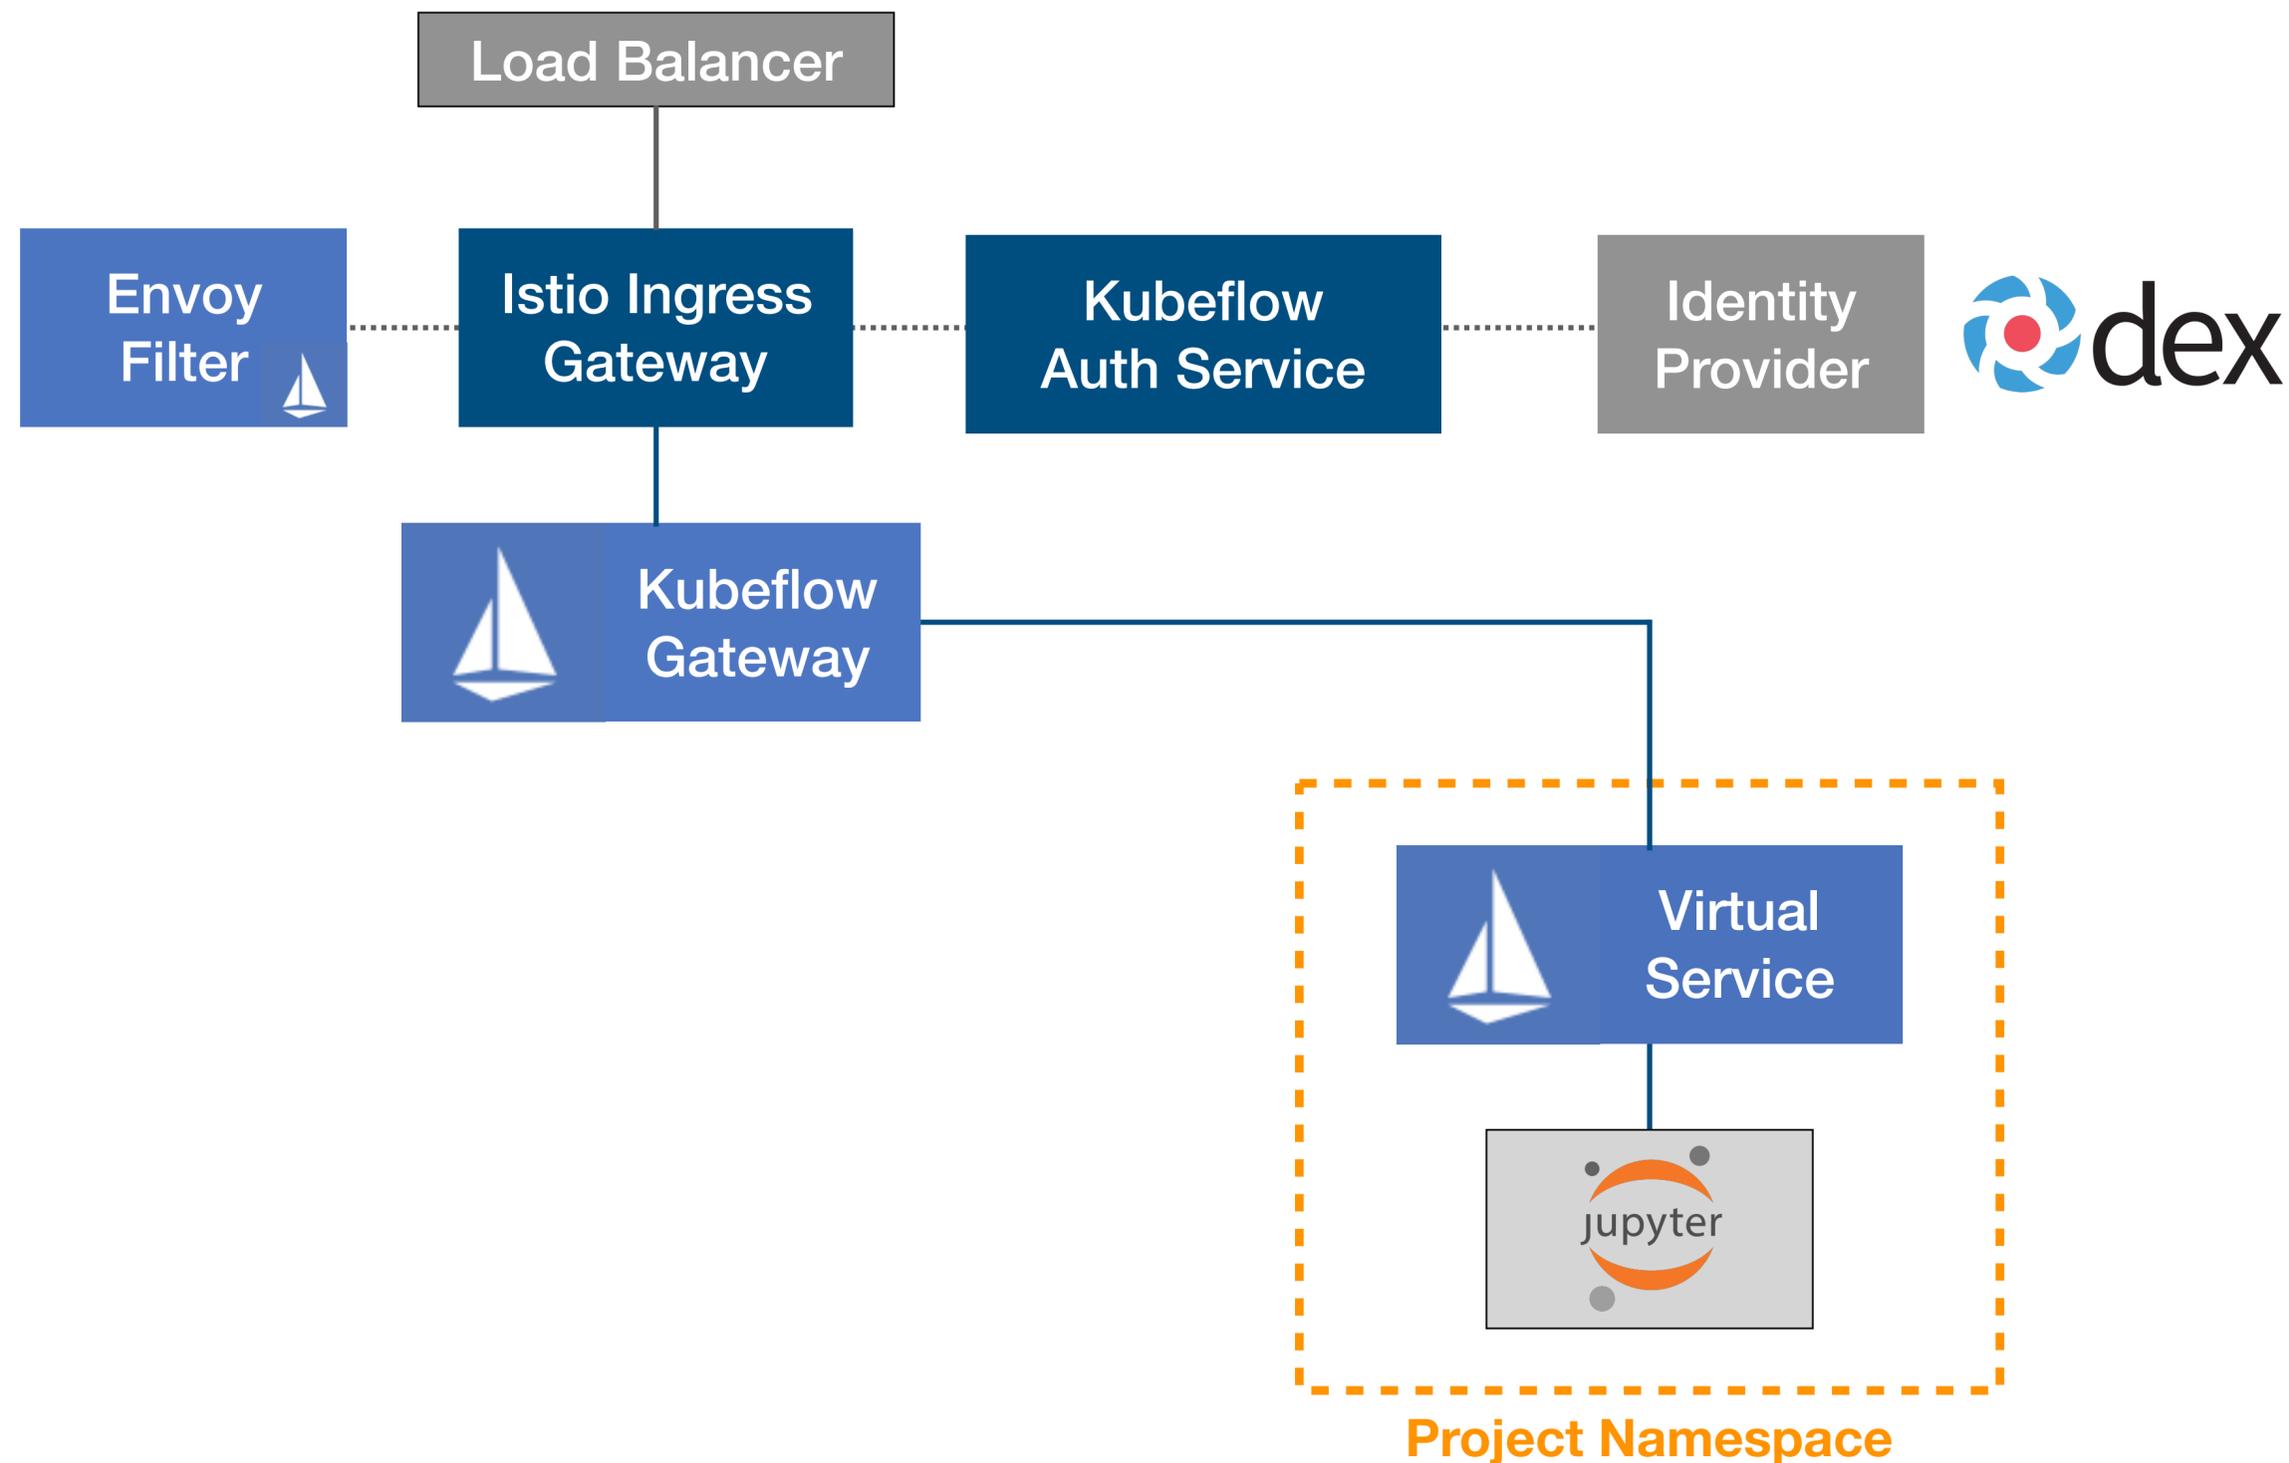


KubeCon



CloudNativeCon

North America 2019



Data and Secrets



KubeCon



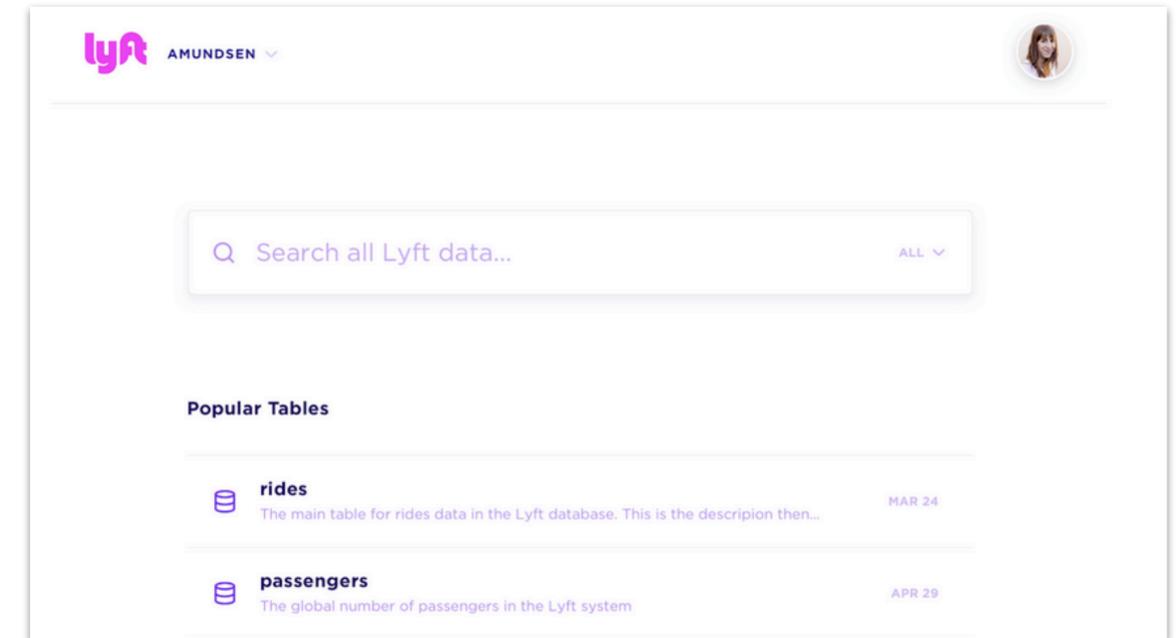
CloudNativeCon

North America 2019

Data

Control access to data via an auditable layer

Allow for data discoverability (with tools like *Amundsen*)



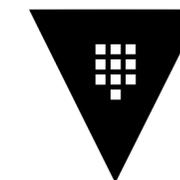
Amundsen

<https://github.com/lyft/amundsen>

Secrets

Use a secret manager

Pods can authenticate to Vault and get secrets loaded in memory



HashiCorp

Vault

Monitoring



KubeCon



CloudNativeCon

North America 2019

Goal:

- Monitor projects for users
- Monitor cluster for MLOps
- Collect any metrics from jobs
- Automated dashboards



Prometheus



Grafana



Lessons learned:

- Grafana Dashboards as *ConfigMaps*
- Allow users to submit dashboards via *GitOps*
- Multiple K8s clusters? Use Federation feature
- Need to scale up prometheus storage? Use *Thanos*

Monitoring cost



KubeCon



CloudNativeCon

North America 2019

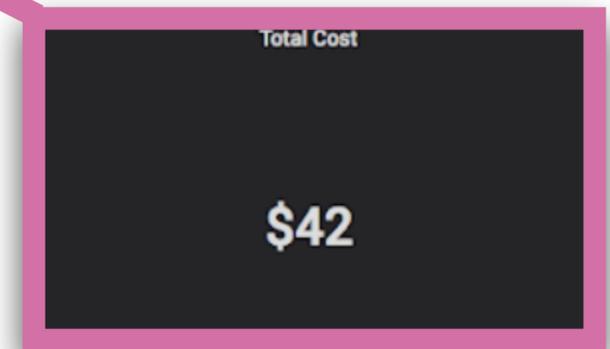
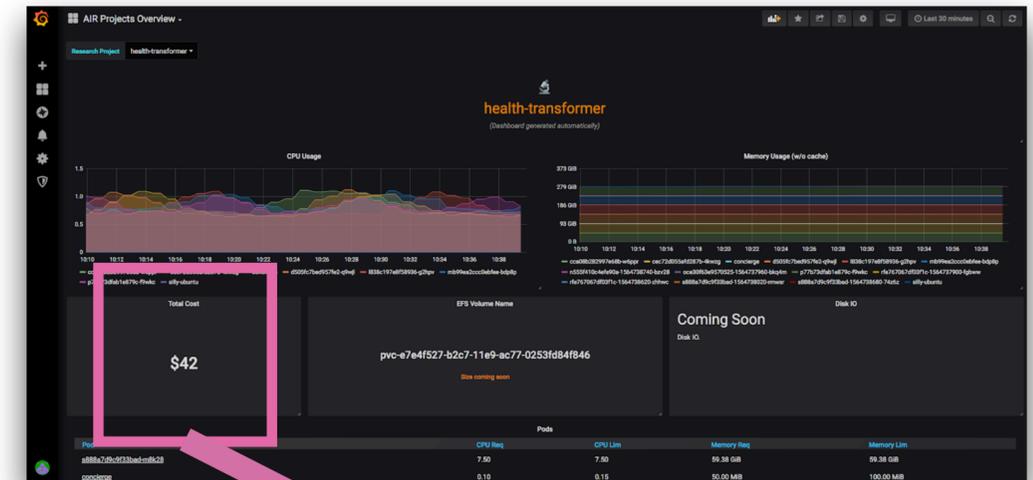
Keeping track of cost in a cloud environment is vital especially in AI/ML

Many proprietary options...

One of them open-sourced their cost model engine:



github.com/kubecost/cost-model



Babylon:



Get
cost per
namespace

Expose
cost as metrics



Pull metrics





KubeCon



CloudNativeCon

North America 2019



Self-service access to ML Toolkit

CLI-based interface

Monitoring

On-demand compute

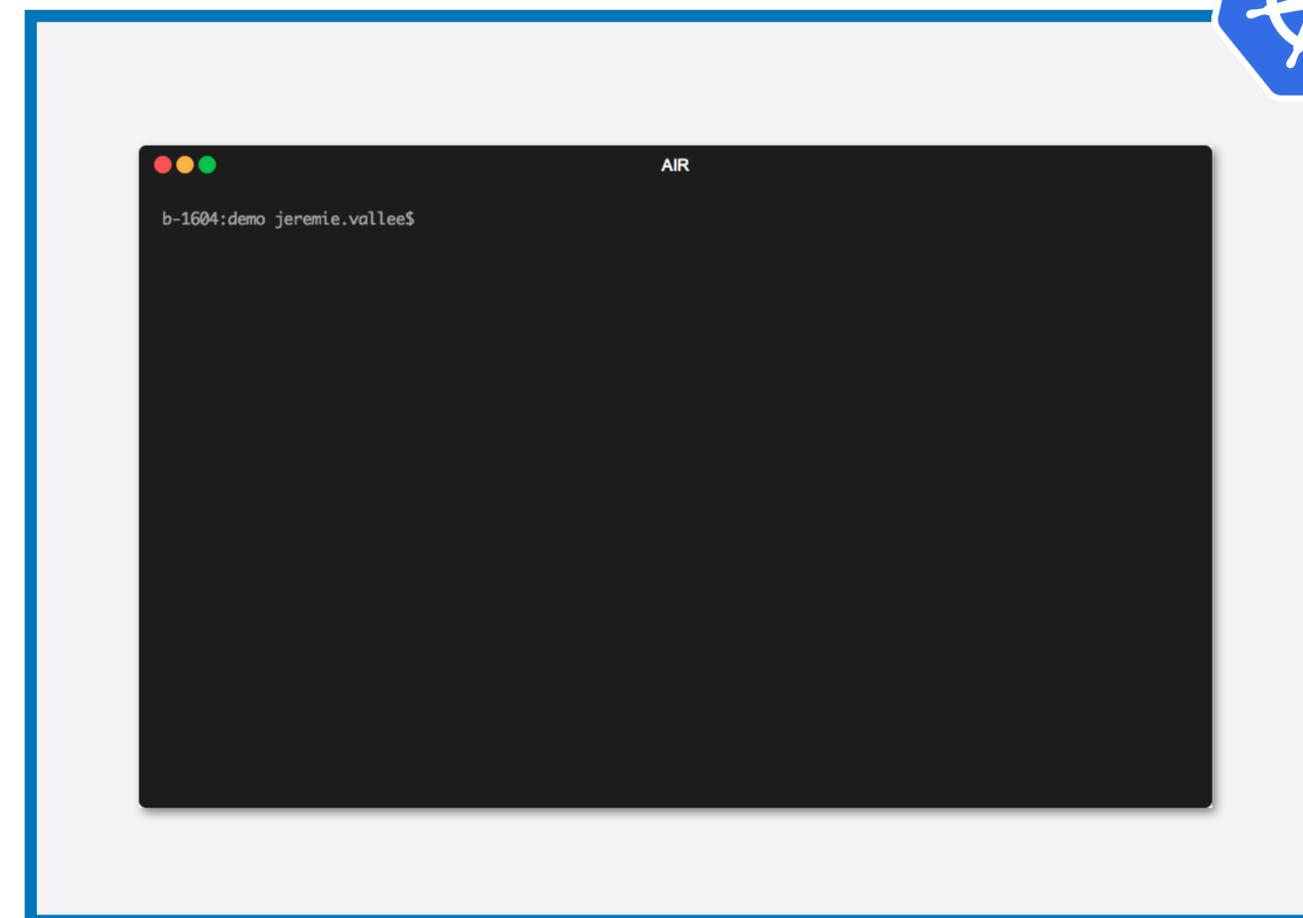
Network RBAC, mTLS



Single cluster

Slow on-boarding (GitOps)

Complex Kubernetes objects



Global platform

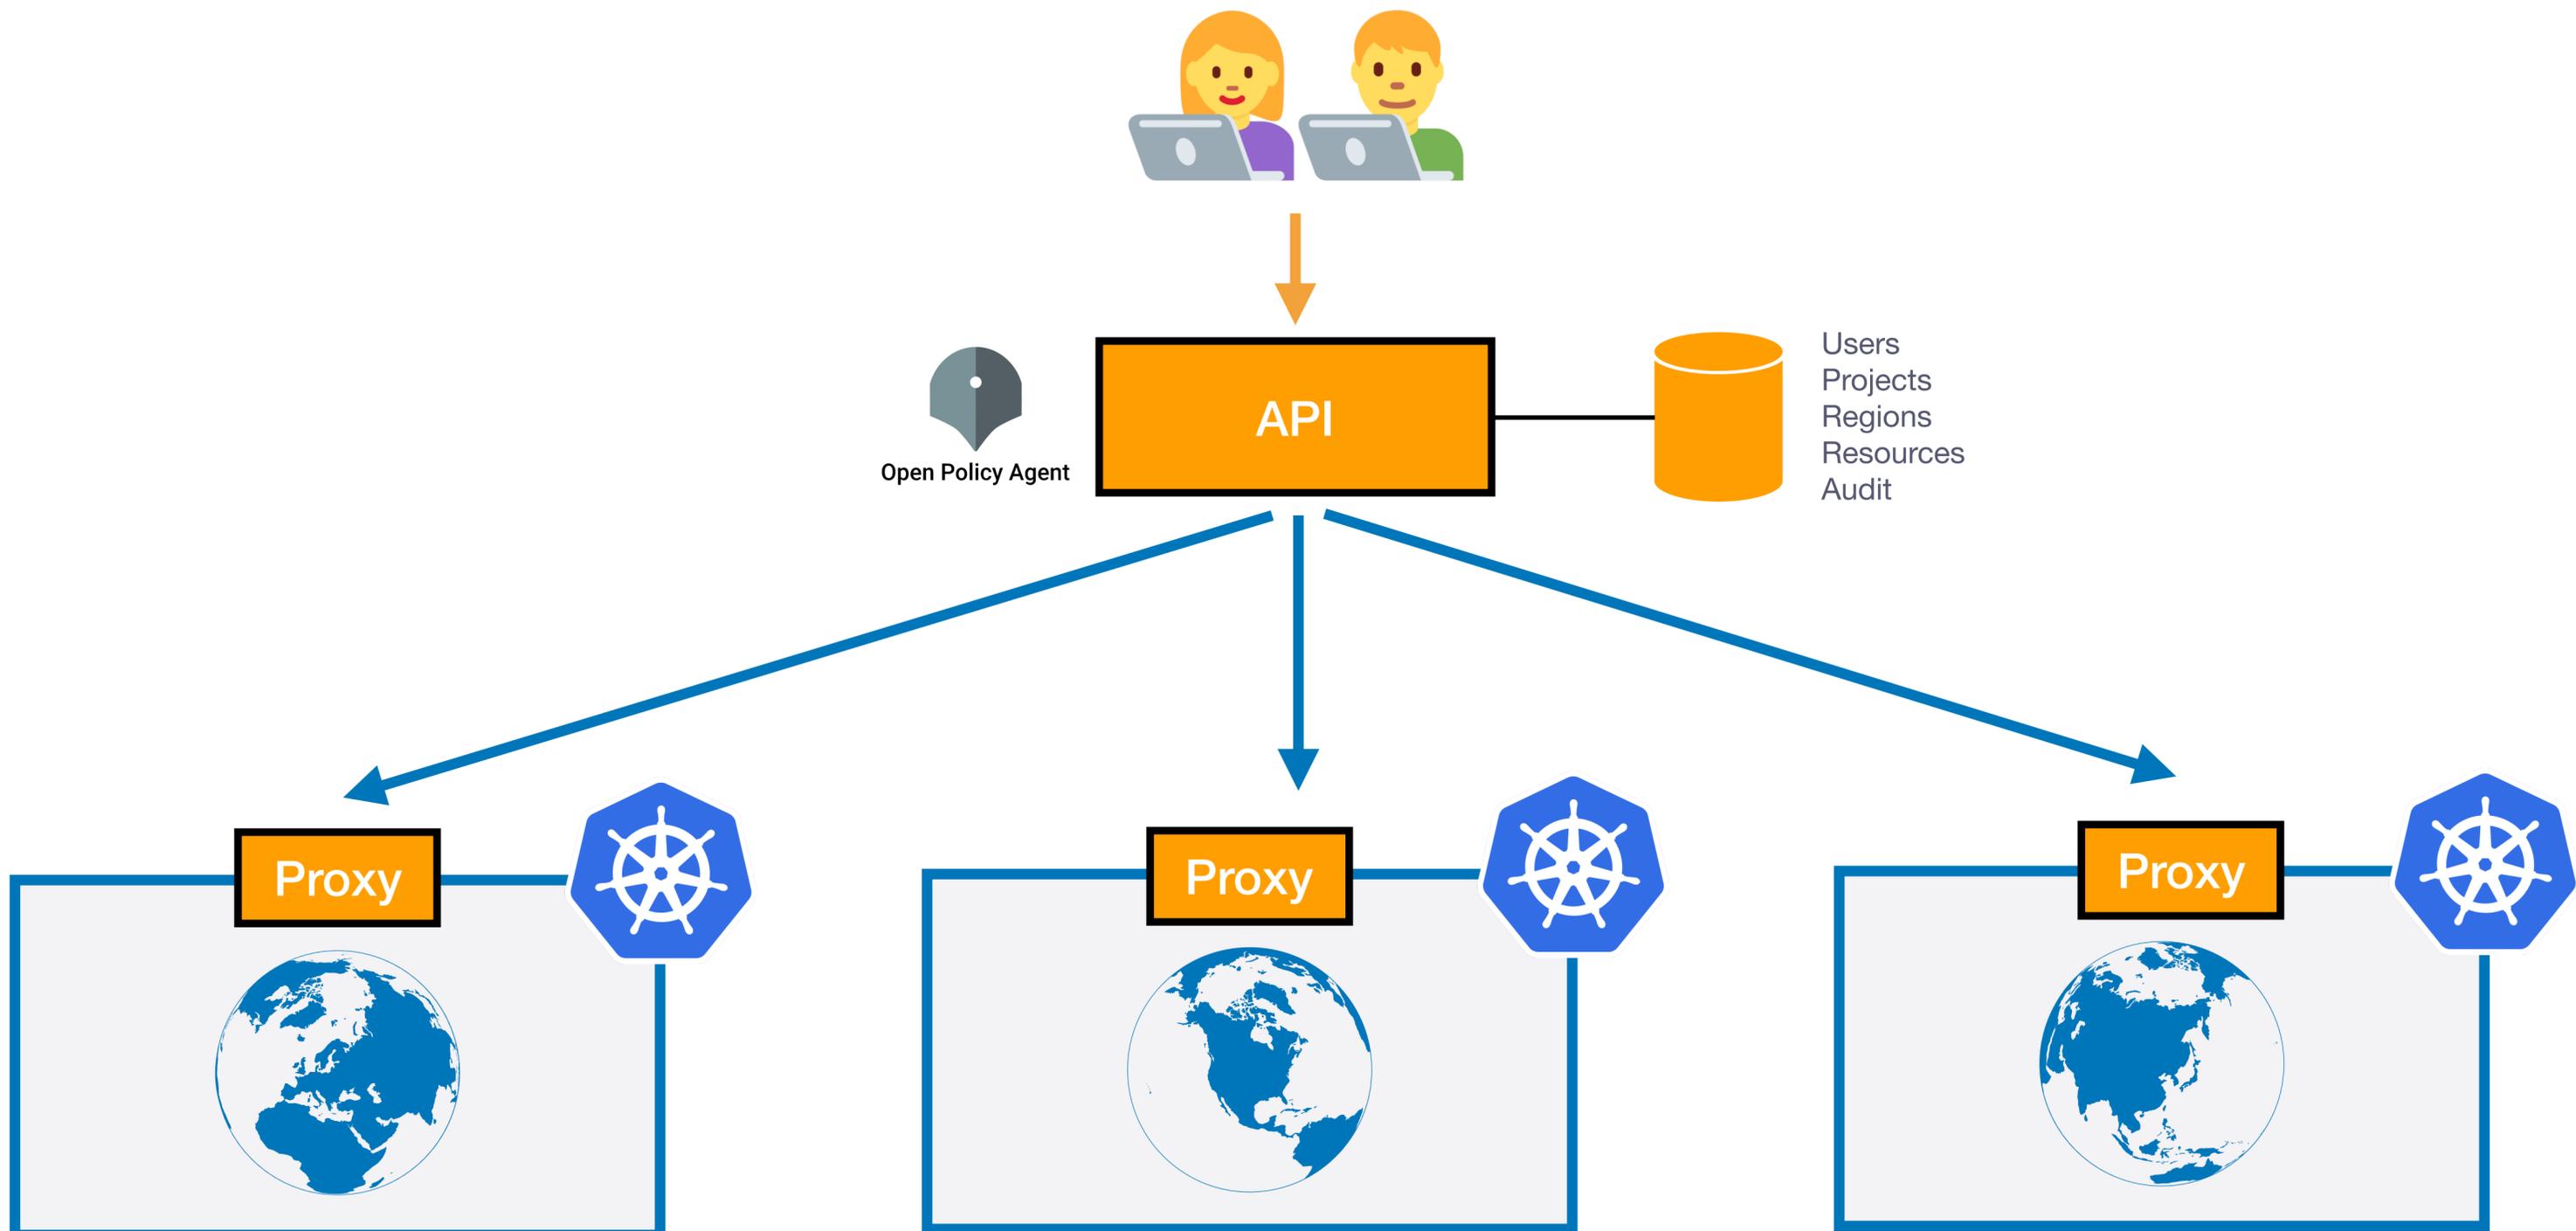


KubeCon



CloudNativeCon

North America 2019



Simplifying resources



KubeCon



CloudNativeCon

North America 2019



```
{
  "kind": "pod",
  "name": "simple-gpu-example",
  "image": "nvidia/cuda:8.0-cudnn5-runtime",
  "command": ["python"],
  "args": ["script.py"],
  "resources": "gpu_medium"
}
```



API



```
---
apiVersion: v1
kind: Pod
metadata:
  name: simple-gpu-example
  namespace: my-project
spec:
  containers:
  - image: nvidia/cuda:8.0-cudnn5-runtime
    command: ["python"]
    args: ["script.py"]
    name: simple-gpu-example
    resources:
      limits:
        memory: "16Gi"
        cpu: "8000m"
        nvidia.com/gpu: 1
    name: "tensorflow"
    volumeMounts:
    - mountPath: /mnt
      name: efs-storage
  restartPolicy: "OnFailure"
  volumes:
  - name: efs-storage
    persistentVolumeClaim:
      claimName: efs
  imagePullSecrets:
  - name: my-deploy-pull-secret
  tolerations:
  - key: "nvidia.com/gpu"
    operator: "Equal"
    value: "true"
    effect: "NoSchedule"
  nodeSelector:
    accelerator: nvidia-tesla-k80
```



KubeCon



CloudNativeCon

North America 2019

Self-service access to ML Toolkit

CLI and/or Web interface

Monitoring

On-demand compute

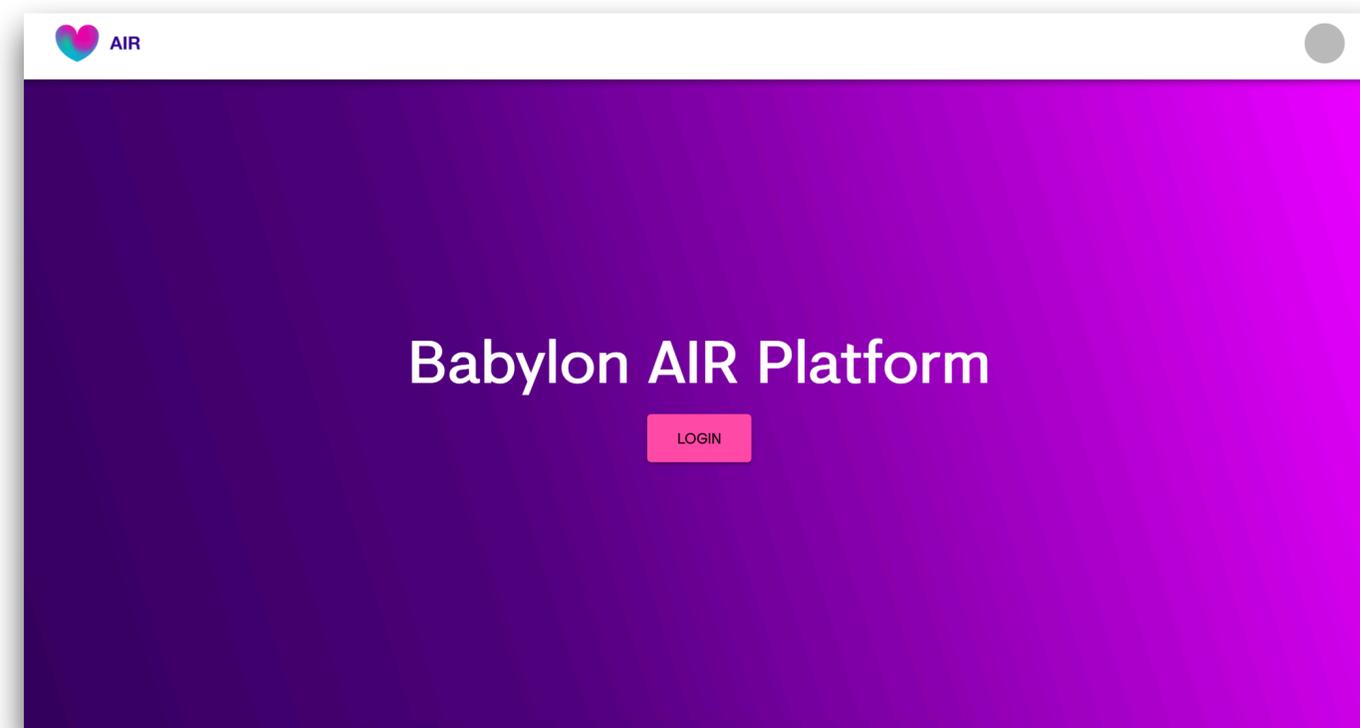
Network RBAC, mTLS

Multi region

Fast on-boarding (UI or API)

Simpler object definitions (heavy-lifting in the backend)

Multi user



Use Case: *Clinical Validation of our Symptom Checker*

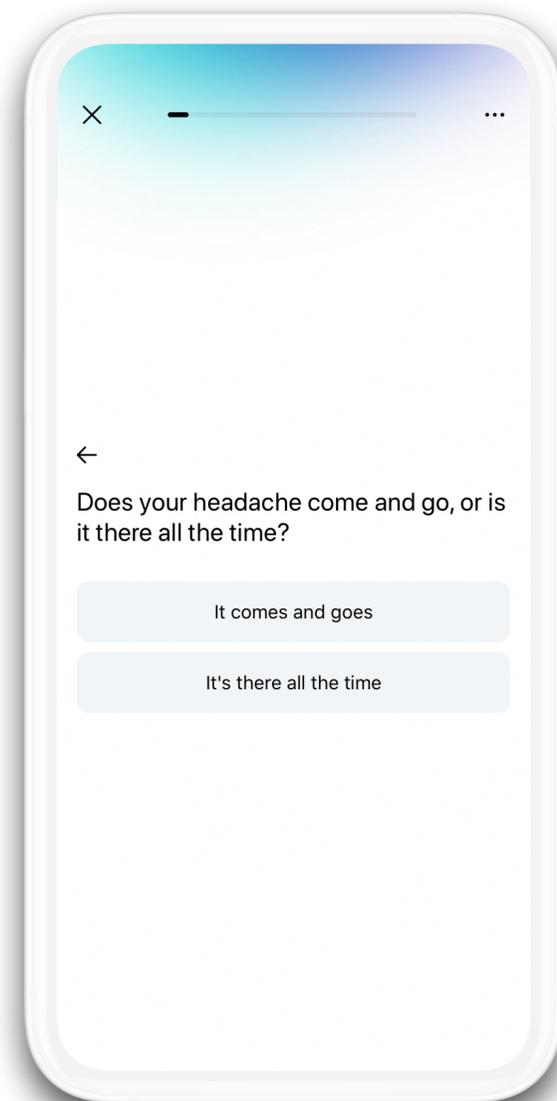


KubeCon



CloudNativeCon

North America 2019



Symptom Checker: Bayesian network

Has two jobs:

1. Get as much relevant evidence from patient as possible
2. Find most likely disease based on evidence received

$$P(D_i | E)$$

Probability of Disease given Evidence

We want to evaluate how well it performs.

Use Case: *Clinical Validation of our Symptom Checker*



KubeCon

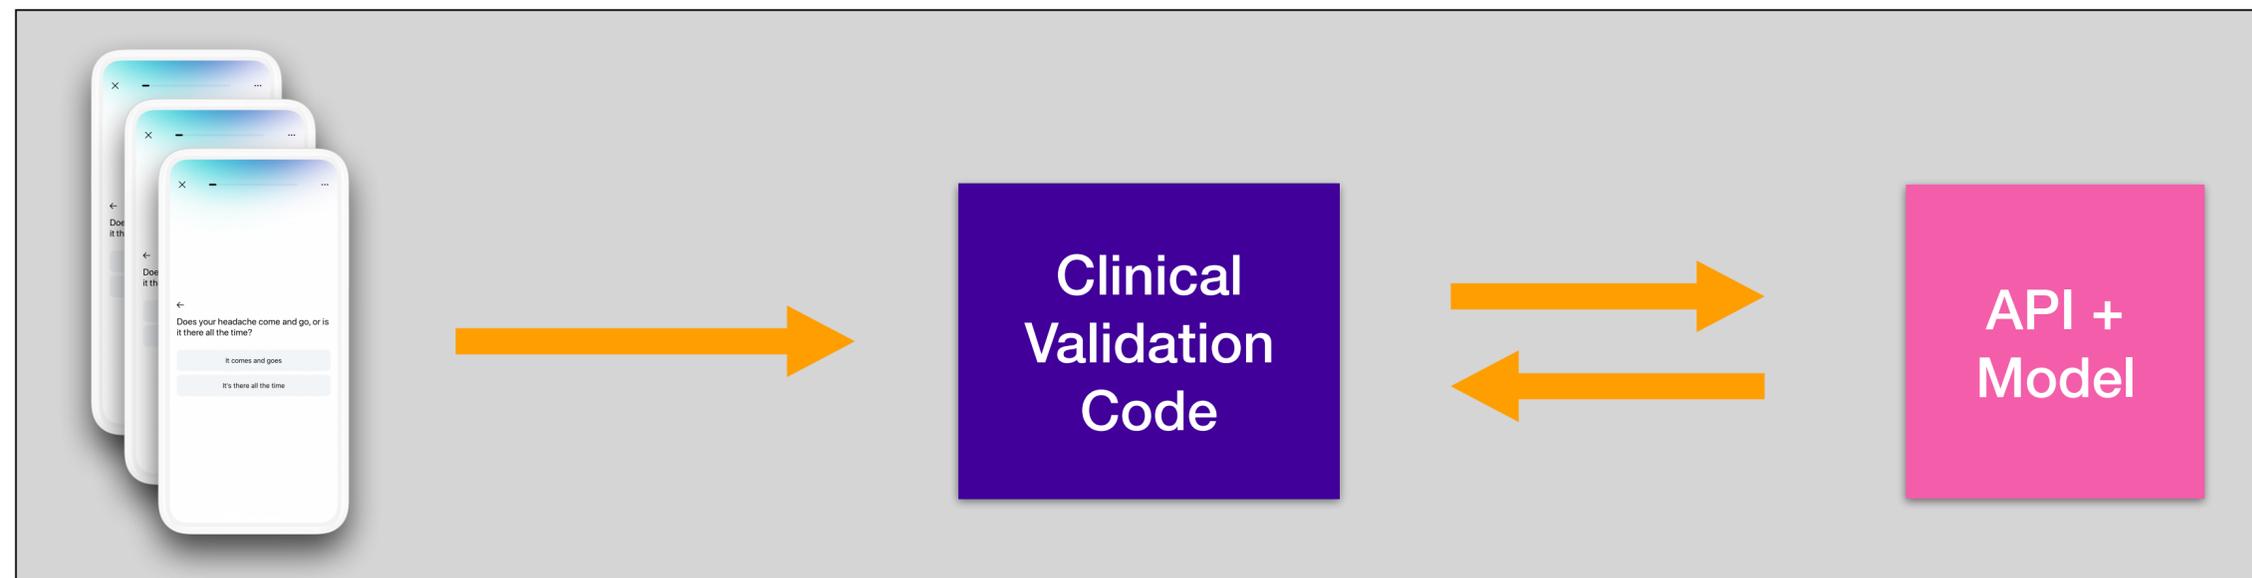


CloudNativeCon

North America 2019

Clinical Validation:

- Simulating patient interaction via use cases generated by doctors
- Evaluating both questions asked by model and outcome
- Original duration: **10 hours** (*and lots of misery*)



*Thousands
of use cases*

*Managing use cases
+
Assessing results*

Serving Model

Use Case: *Clinical Validation of our Symptom Checker*

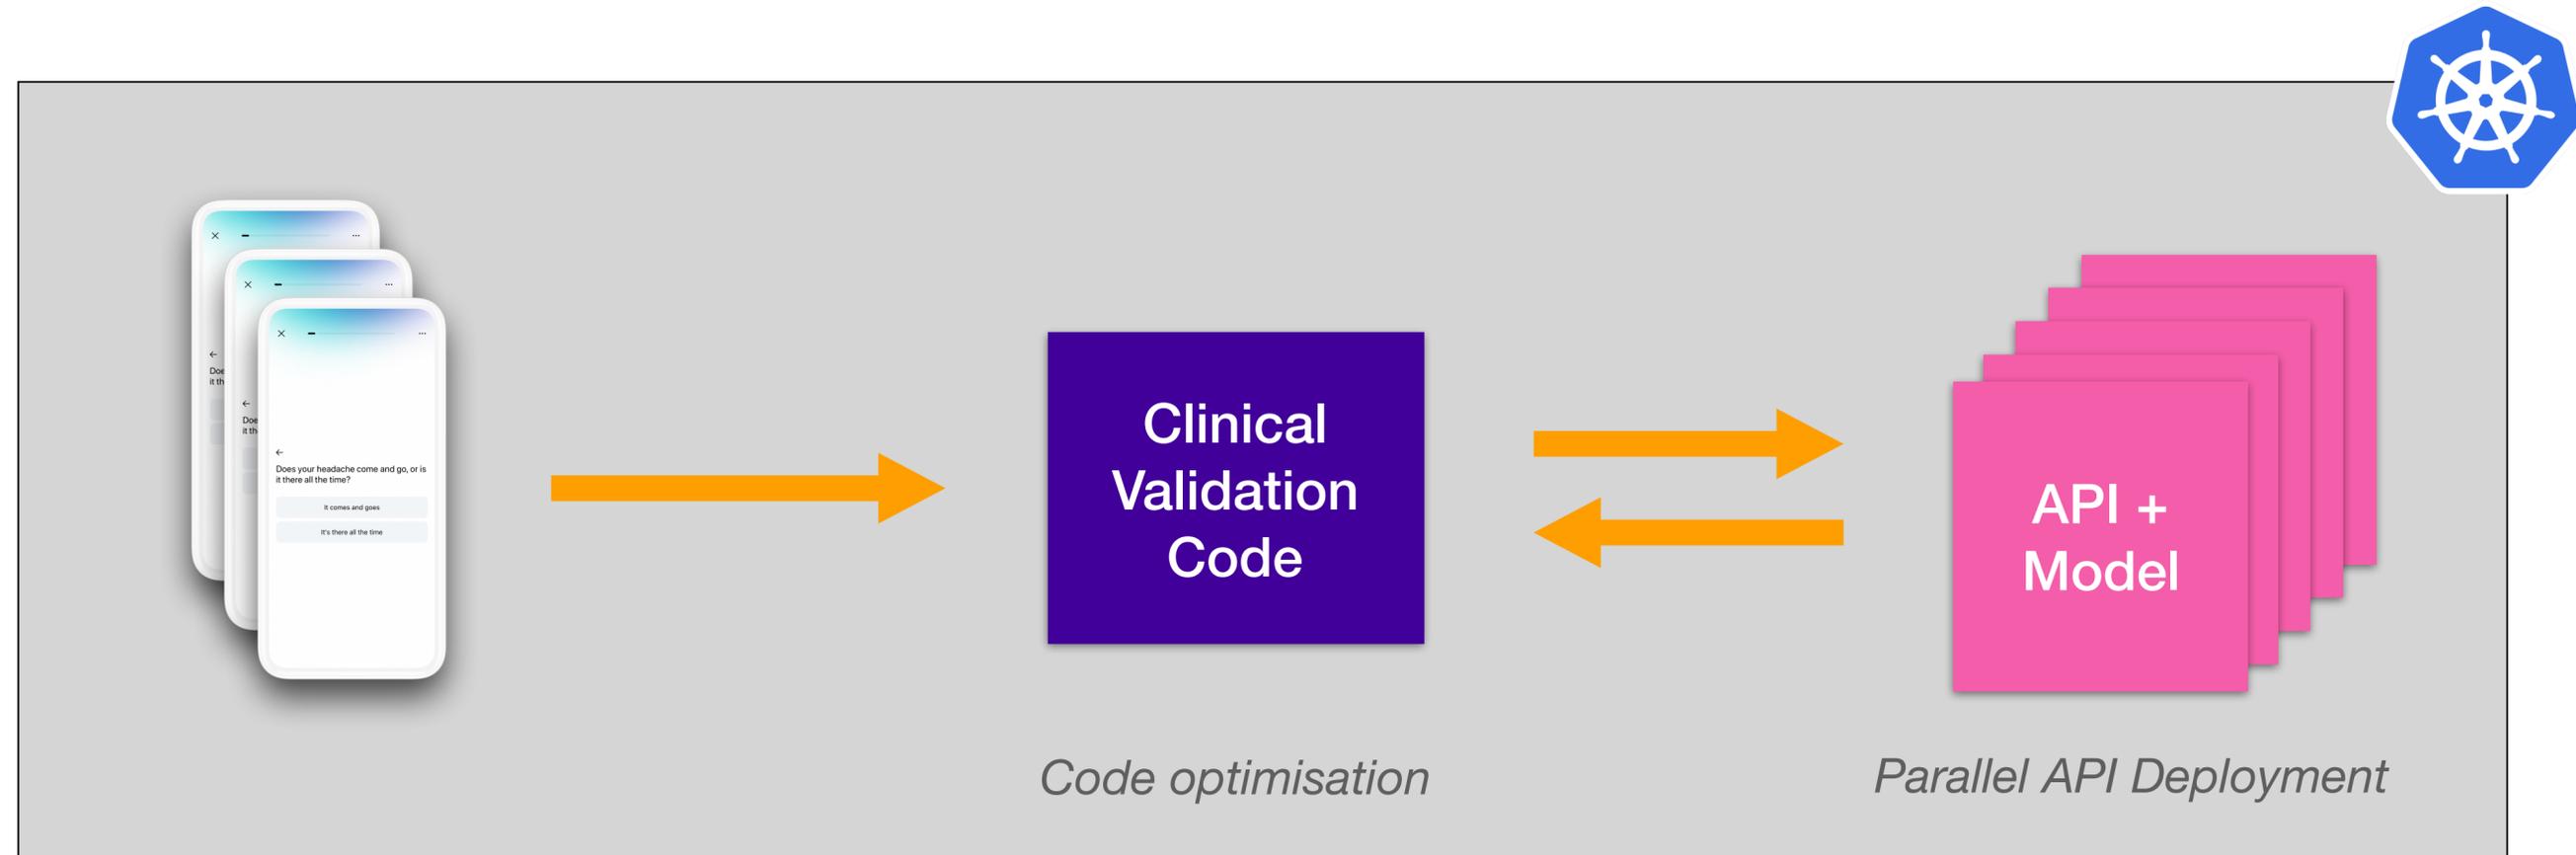


KubeCon



CloudNativeCon

North America 2019



New duration: **< 20 minutes**

Now running on every Pull Request

Improving feedback loop → faster iterations → **increasing safety and quality of models**

Next steps



KubeCon



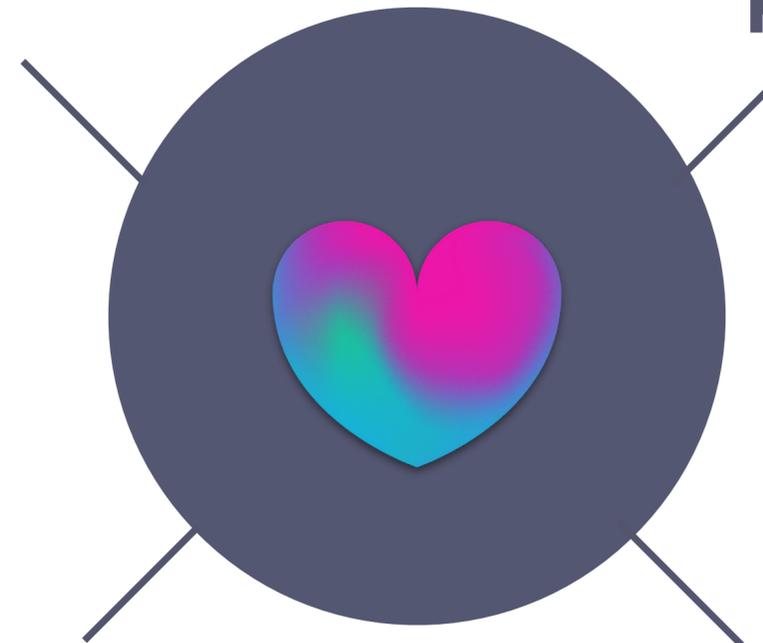
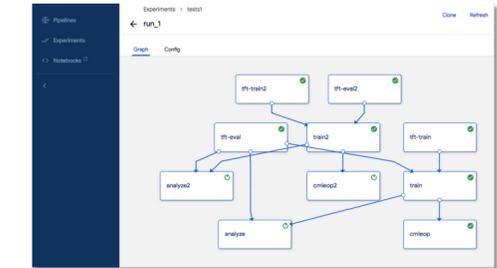
CloudNativeCon

North America 2019



Better serving

Integrating Kubeflow Pipelines

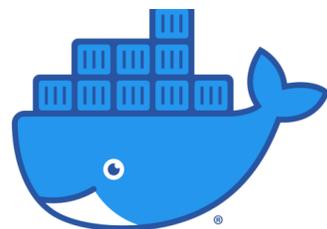


Improving user experience

Better Metadata Tracking



Sacred



Wrapping up



KubeCon



CloudNativeCon

North America 2019

Kubeflow:

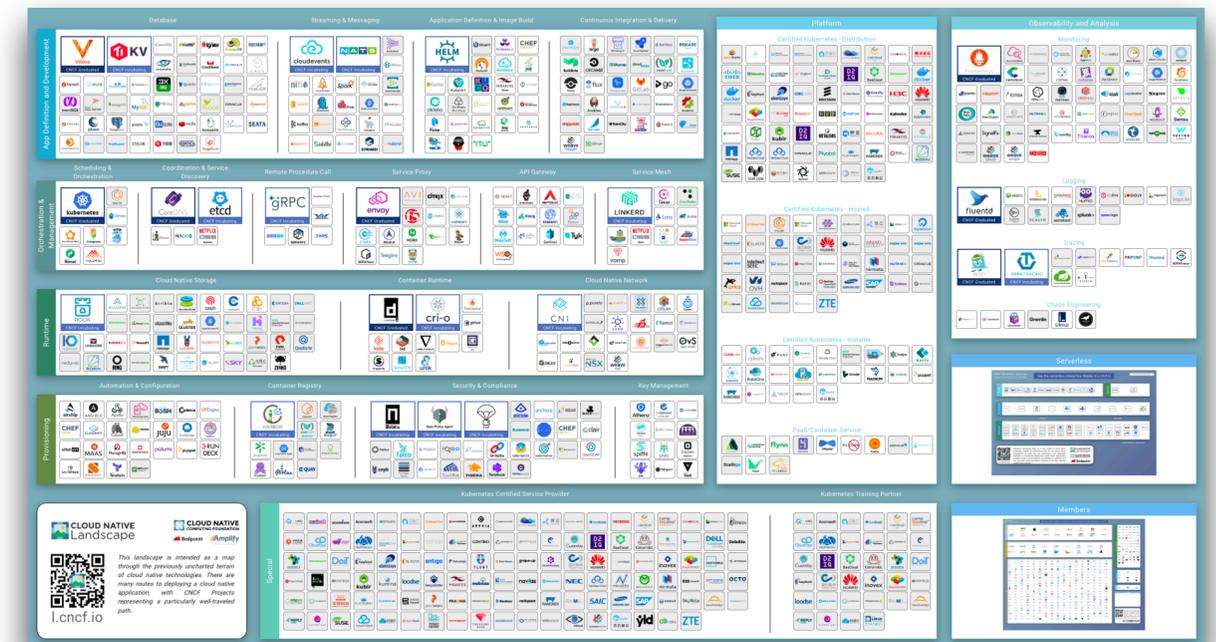
- Most complete ML toolkit for Kubernetes
- Great modularity
- Easy to get started

Security and compliance on K8s:

- Many open-source tools out there can help
- Have a look at the Cloud Native landscape and start from there

MLOps:

- Enabling AI/ML teams with tooling and infrastructure
- Always ask: *what are the big pain points for your AI/ML teams?*
- Focus on 1 pain point, build proof-of-concept, then add as feature

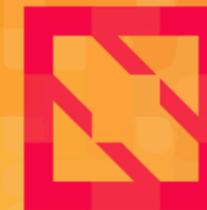


<https://landscape.cncf.io/>

Thank you!



KubeCon



CloudNativeCon

North America 2019



@jeremievallee



jeremie-vallee

*We're hiring in **UK** and **USA!***

