



K8s, What's Next?

Compute Architectures, Service Mesh, API Managers, Kubeflow
& FaaS on K8s

Pete Johnson
PSE, GPO

Agenda

- Compute Architectures
- Service Meshes + API Gateways
- Kubeflow
- FaaS on K8s

Compute Architectures

Two Prominent AWS re:Invent Announcements

Introducing Firecracker, a New Virtualization Technology and Open Source Project for Running Multi-Tenant Container Workloads

Posted On: Nov 26, 2018

<https://aws.amazon.com/about-aws/whats-new/2018/11/firecracker-lightweight-virtualization-for-serverless-computing/>

Announcing AWS Outposts

Posted On: Nov 28, 2018

<https://aws.amazon.com/about-aws/whats-new/2018/11/announcing-aws-outposts/>


AWS Stack



Nitro Background

- Architecture behind latest instance types, including bare metal
- Swappable CPUs
 - Intel, AMD, Nvidia, ARM

Nitro: One year later



AWS Nitro

Launched in November 2017

In development since 2013

All new launches use Nitro

Purpose-built hardware/software

Hypervisor built for AWS

AWS
Big
Bet #1

AWS
re:Invent

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Nitro Architecture

Nitro in three parts

Nitro Cards



VPC Networking
Amazon Elastic Block Store
(Amazon EBS)
Instance Storage
System Controller

Nitro Security Chip



Integrated into motherboard
Protects hardware resources
Hardware Root of Trust

Nitro Hypervisor



Lightweight hypervisor
Memory and CPU allocation
Bare Metal-like performance

aws
re:Invent

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Nitro Cards

ENA PCIe Controller



VPC Data Plane

NVMe PCIe Controller



Transparent Encryption

NVMe PCIe Controller



EBS Data Plane

System Control



Root of Trust

aws
re:Invent

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Firecracker: What's a Micro-VM?

Firecracker Micro-VMs




Security

Minimal device model reduces memory footprint and attack surface area



Speed by design

User-space code in <125ms, 150 microVM per second per host



Scale and efficiency

Low memory overhead with high density (thousands) of microVMs on each server

Serverless
@goserless

Following

AWS just did a live #Firecracker test where they spun up 4000 microVMs. The longest one took 219 ms 😲
#reinvent2018 #reinvent

10:43 AM - 27 Nov 2018

<https://twitter.com/goserless/status/1067489201808130049>

- Used under the hood for Lambda since 2014
- Could be used for 1 VM per Container EKS/ECS (like OpenStack Kata)
- Micro-VM sizes 1CPU 128Mb to 3GB (based on Lambda)

Outposts: Sometime in 2019 . . .

- 80” cabinet and smaller sizes shipped to customer
- Control plane stays in AWS AZ
- Expected to offer EC2 and EKS, *could* offer Firecracker and Lambda
- Priced similar to reserved instances, but with equipment to return

Nitro: Anywhere you need it



AWS Outposts

Nitro hardware and software in your data center

Access via standard AWS API and console

Deploy apps to Outposts using AWS services

AWS
re:Invent

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Extension of AWS Region

AWS Region
(N. California)

AWS Outposts
Sacramento



AWS Outposts
Los Angeles



AWS Outposts
Houston



AWS
Big
Bet #2:
Latency

Compute Architectures Summary

- Firecracker Mini-VMs could disrupt how containers run in virtualization
 - Similar to OpenStack Kata containers

- Outposts could disrupt where containers get hosted on hardware
 - But how will it deal with latency between masters and nodes?

Service Mesh & API Gateways

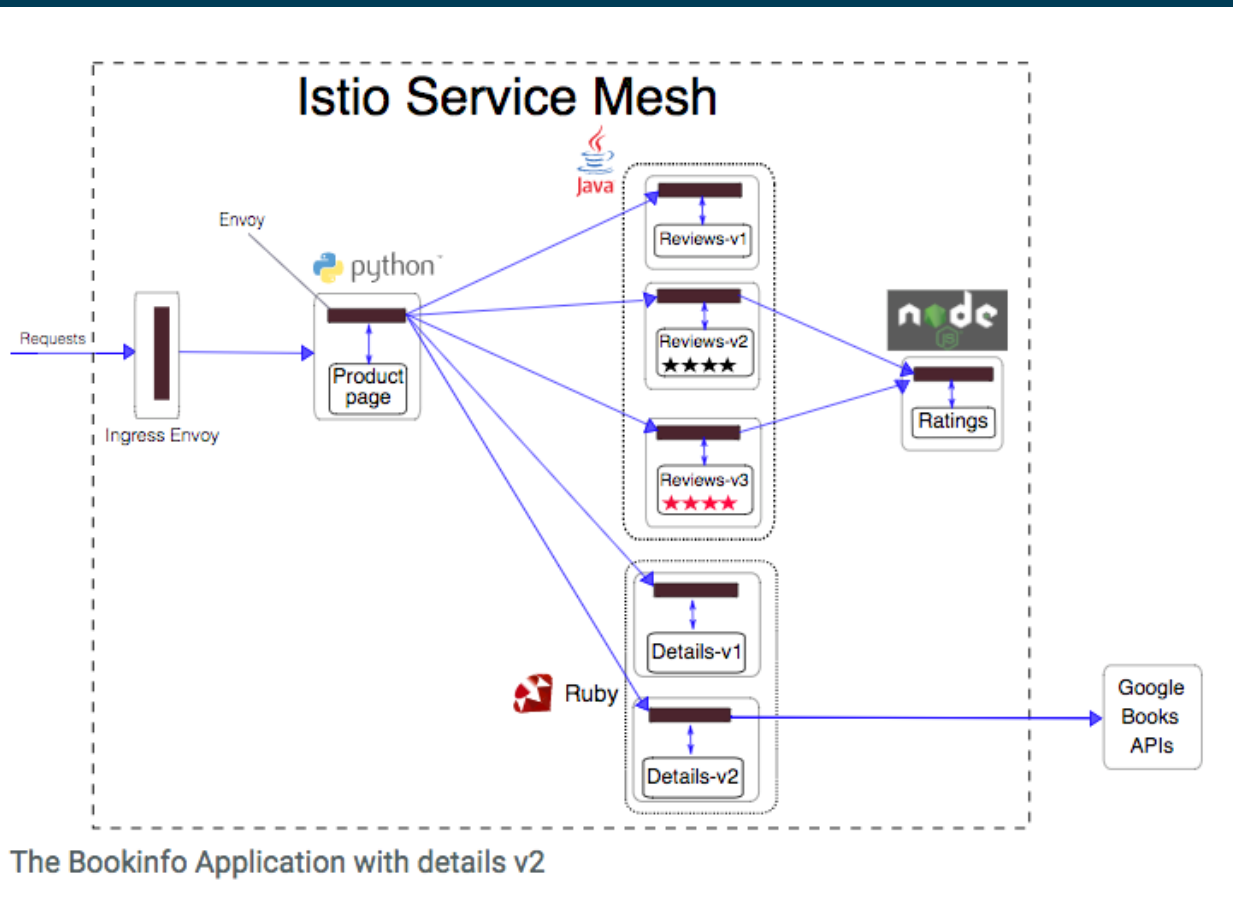
What's a Service Mesh?



What Is a Service Mesh?

- Infrastructure layer for service to service communication
- A mesh of proxies injected as sidecars that support numerous formats with intelligent routing rules between endpoints
 - Further example: Contiv VPP is a sidecar
- *Can* inspect API transactions at Layer 7 or Layer 3/4

Application Example

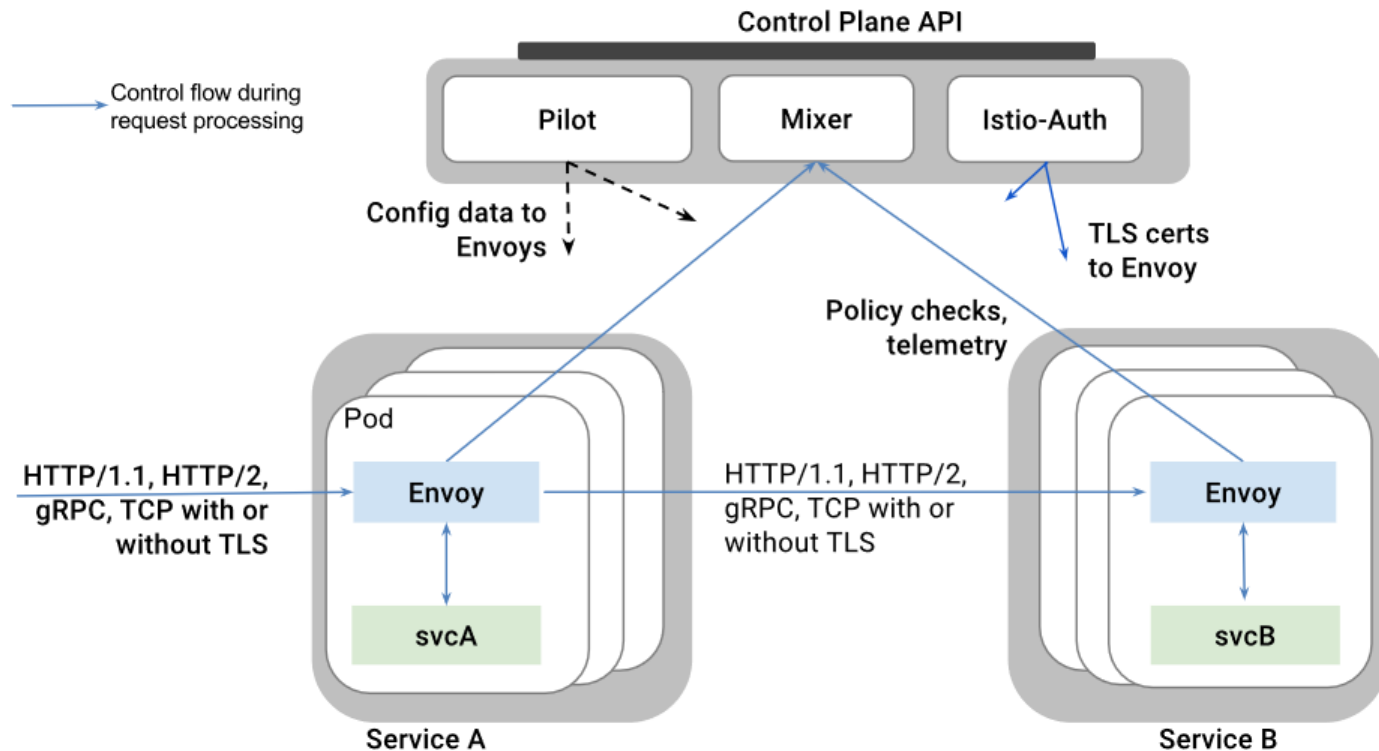


Core Istio Architectural Components

- Pilot (istio-pilot)
 - Handles service discovery and config data
 - Provides the Envoy proxies with the mesh topology and route rules
- Mixer (split into telemetry function and policy function)
 - Istio-telemetry collects telemetry data and passes to back-end systems
 - Istio-policy provides more sophisticated policy checks than possible with pilot route rules alone.
 - Easily pluggable
- Envoy
 - A proxy attached to every microservice
 - The connection point for a microservice to attach to the mesh

Istio Architecture

(<https://istio.io/docs/concepts/what-is-istio/overview.html>)

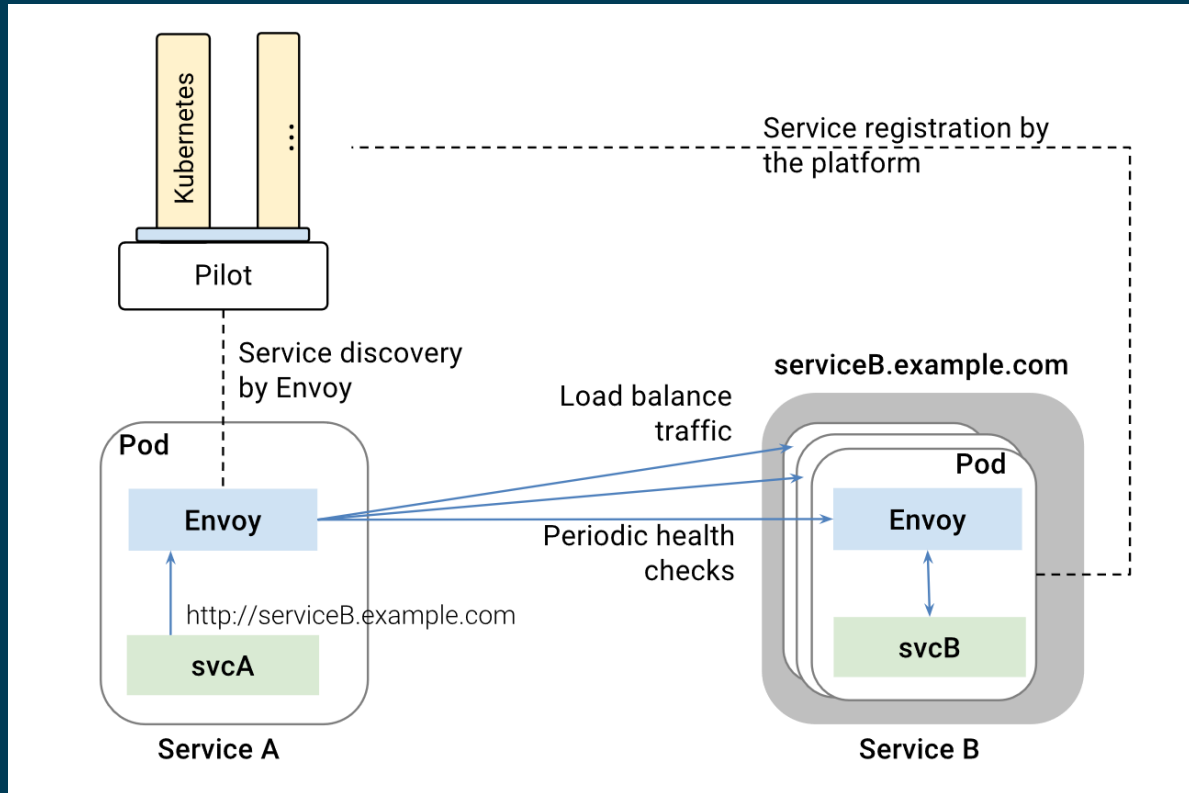


Istio: Why Would I Want To Use It?

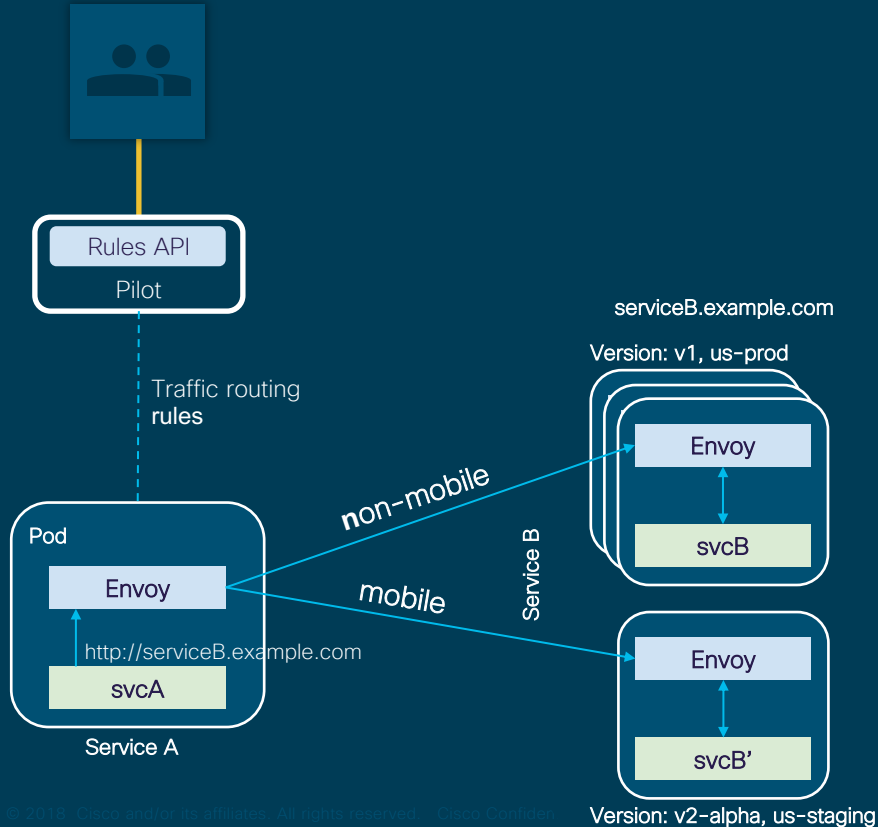
- Eases the building and management of applications comprised of loosely-coupled distributed microservices
- Application visibility for operators
- Frees app developers from developing custom telemetry, authentication, authorization and networking infrastructure
- Policy driven operations
- Tighter application security and fault tolerance
- Rich set of layers 3/4/7 traffic routing and load balancing capabilities

Discovery & Load Balance

(<https://istio.io/docs/concepts/traffic-management/load-balancing.html>)



Service Routing



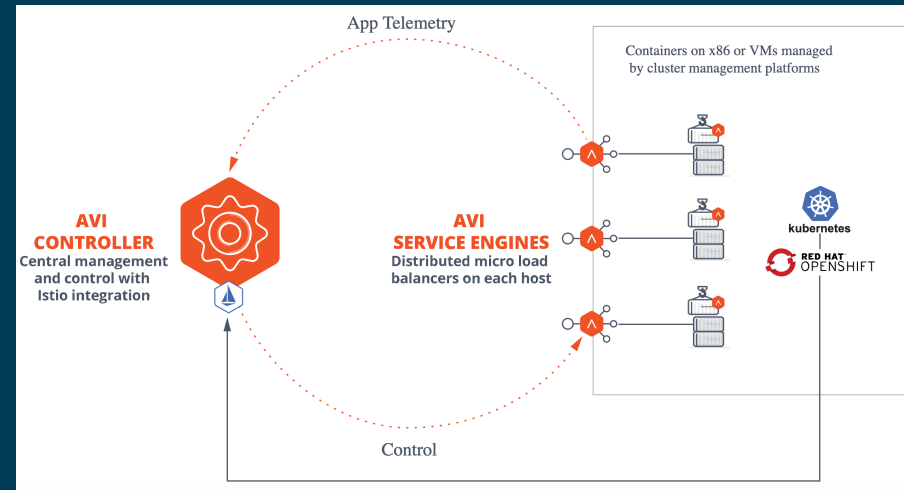
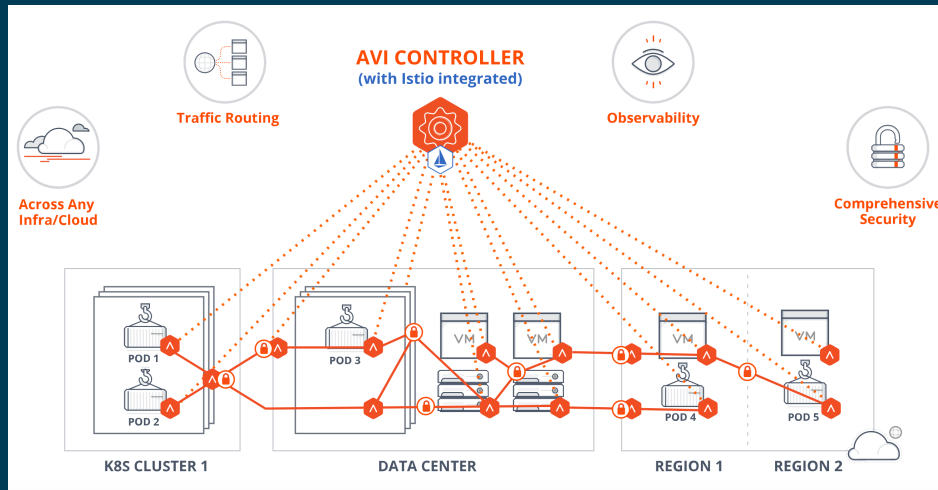
Default rule

```
apiVersion: config.istio.io/v1alpha2
kind: RouteRule
metadata:
  name: svcB-default
  namespace: default
spec:
  destination:
    name: serviceB
  precedence: 1
  route:
  - labels:
    version: v1
```

Mobile rule

```
apiVersion: config.istio.io/v1alpha2
kind: RouteRule
metadata:
  name: svcB-mobile
  namespace: default
spec:
  destination:
    name: serviceB
  precedence: 2
  match:
    request:
      headers:
        user-agent:
          exact: mobile
  route:
  - labels:
    version: v2-alpha
```

AVI Controller with Istio



- Application Maps
- Log Analytics

- Client Analytics
- Security Analytics

Why Use an API Gateway?

- API Time to Market
- Usage Metrics
- Security
 - Throttling
 - API Keys
- Monetization

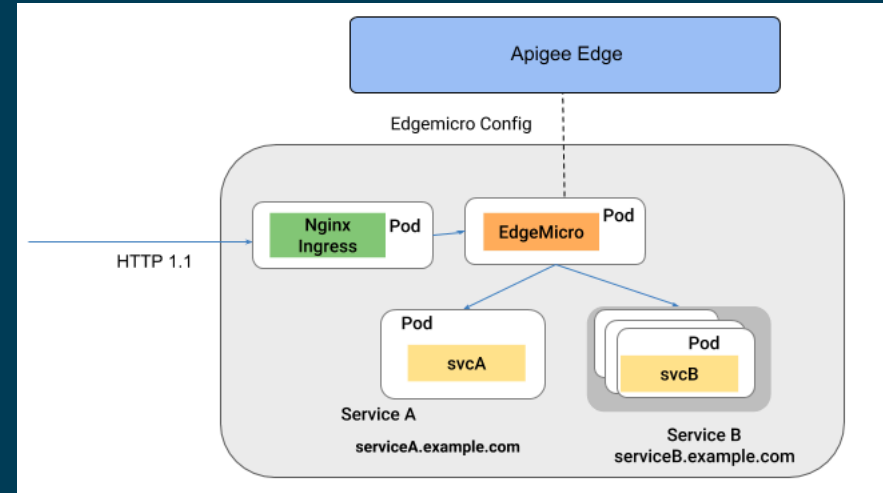
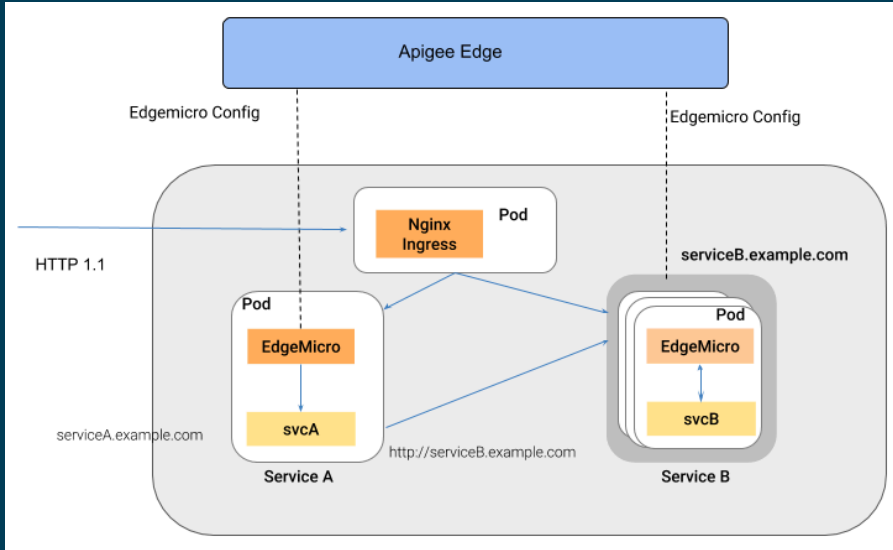


What is APIGee?



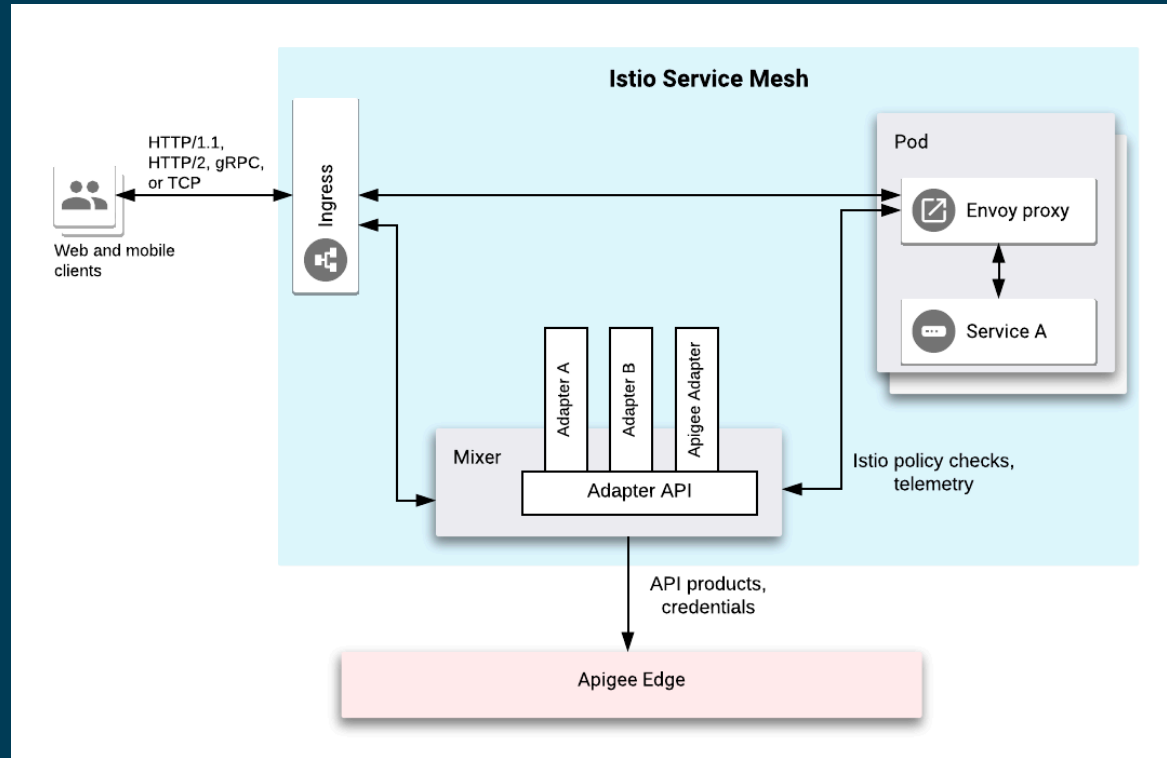
- API Gateway
- Acquired by Google in September of 2016
- Part of the Cisco Hybrid Cloud on Google Cloud solution
 - As a front end to legacy systems that need modern REST APIs
 - K8s integrations on the following slides came after our partnership

APIGee Edge Configuration Examples (Non-Istio)



<https://docs.apigee.com/api-platform/integrations/kubernetes/k8s-introduction>

APIGee Edge and Istio



Service Mesh and API Gateways Summary

- Service Meshes ease service communications within a set of microservices
 - Take on properties of DNS, load balancing, health checks, and other concepts familiar to networking, but for services
 - Istio is the primary service mesh on the market today
- API Gateways can be used to expose REST APIs in a secure, meterable way
 - Starting to build on top of Istio for microservices-based back ends
 - APIGee is a popular commercial API Gateway

Kubeflow

Cisco Contribution to Kubeflow

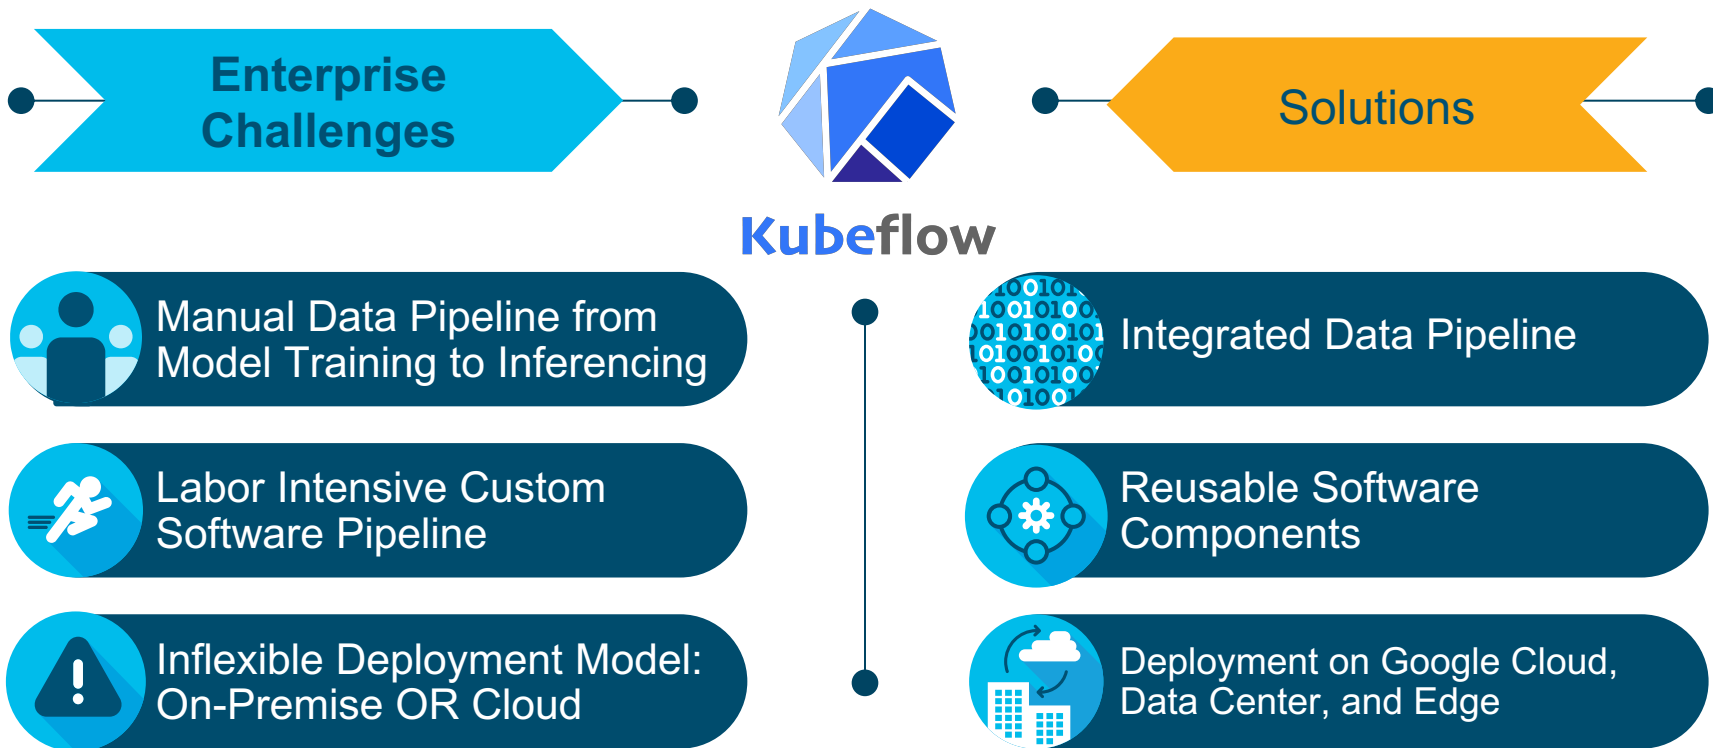
Over 2.8M Lines of Code with 3 Major Proposals



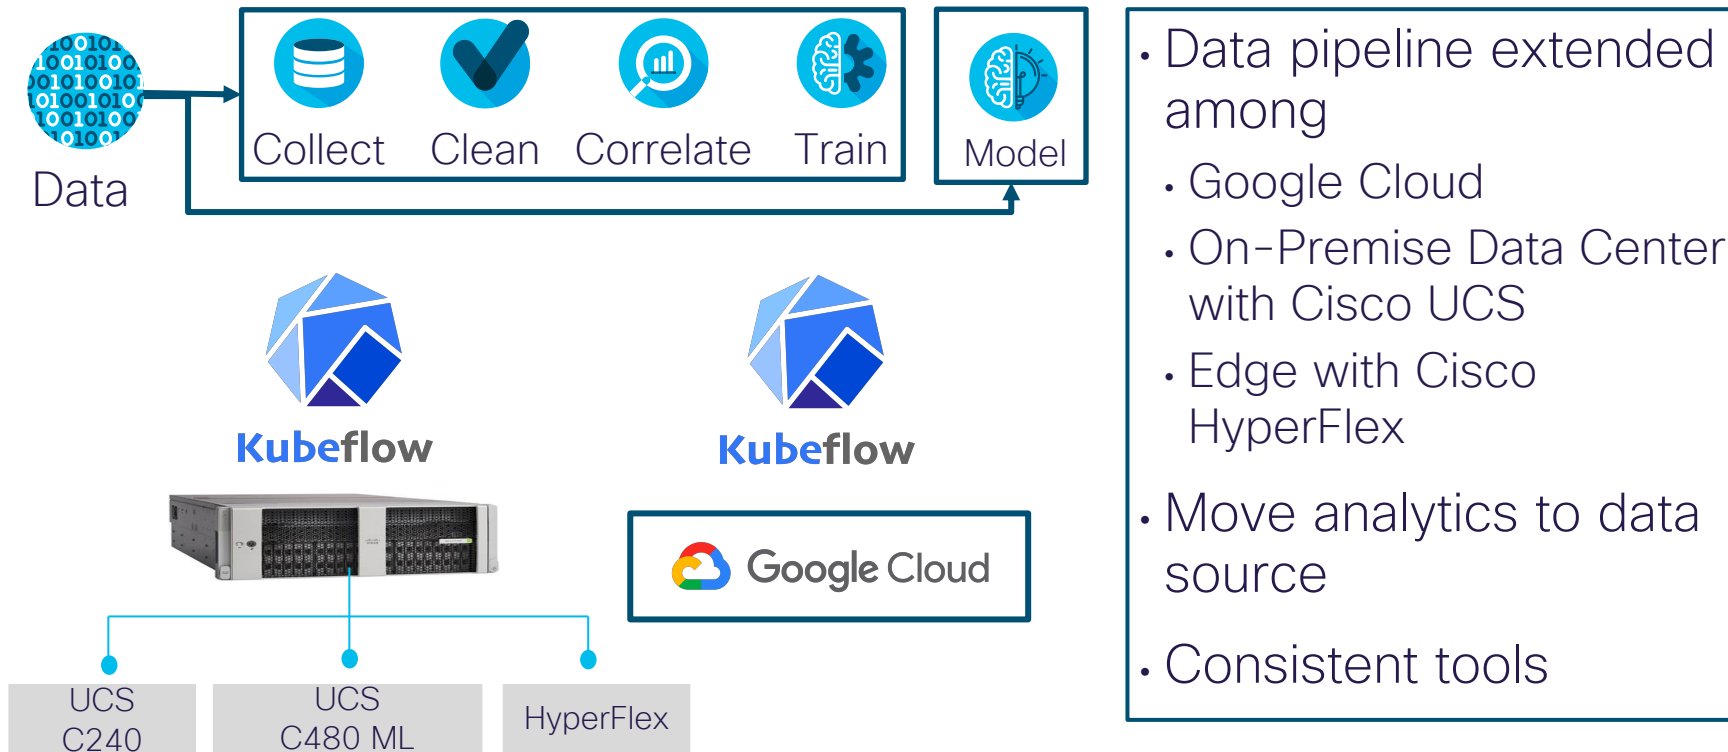
- Thought leadership to expand Kubeflow charter to include hybrid cloud (#ConsistentAI)
- Kubebench: Originated and implemented benchmark for Kubeflow implementation
- PyTorch Operator: Continuous improvement and maintenance

- Katib:
 - Hyperparameter search
 - AutoML with Neural Architectural Search
- Improve on-premise user experience
- Two out top 5 contributors were from Cisco for version 0.3

Challenges of Deploying Data Pipeline



Kubeflow Pipeline on Google Cloud and Cisco UCS and HyperFlex

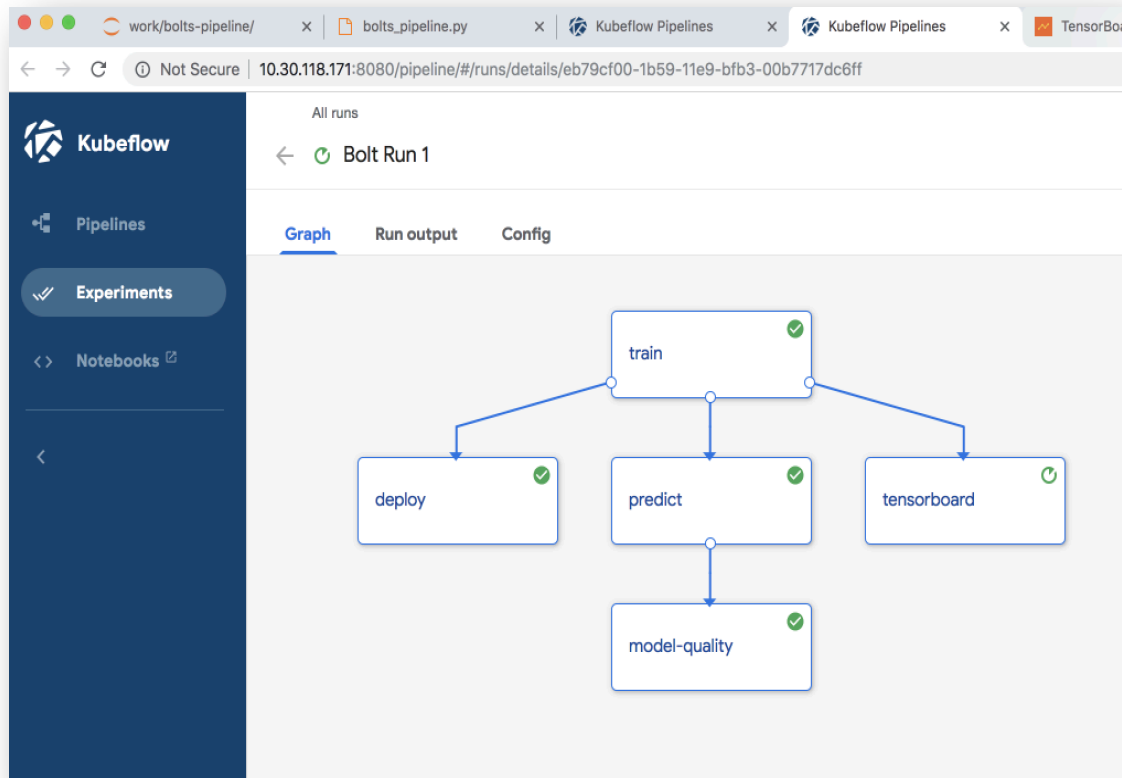


Joint Use Case: Consistent ML Tools on Google Cloud and Cisco UCS

Problem: How to improve selection of bolts of right size?



VS.



Google and Cisco Kubeflow Partnership

- Deploy data pipeline from edge to data center to cloud with Cisco HyperFlex, UCS, and Google Cloud
- Actionable insight wherever data is located
- Consistent machine learning and data tool chain
- Cisco contributing to multiple aspects of Kubeflow, including
 - Inclusion of hybrid cloud use case, Katib, Kubebench, PyTorch Operator, AutoML/Neural Architectural Search and On-premise support.

Kubeflow Summary

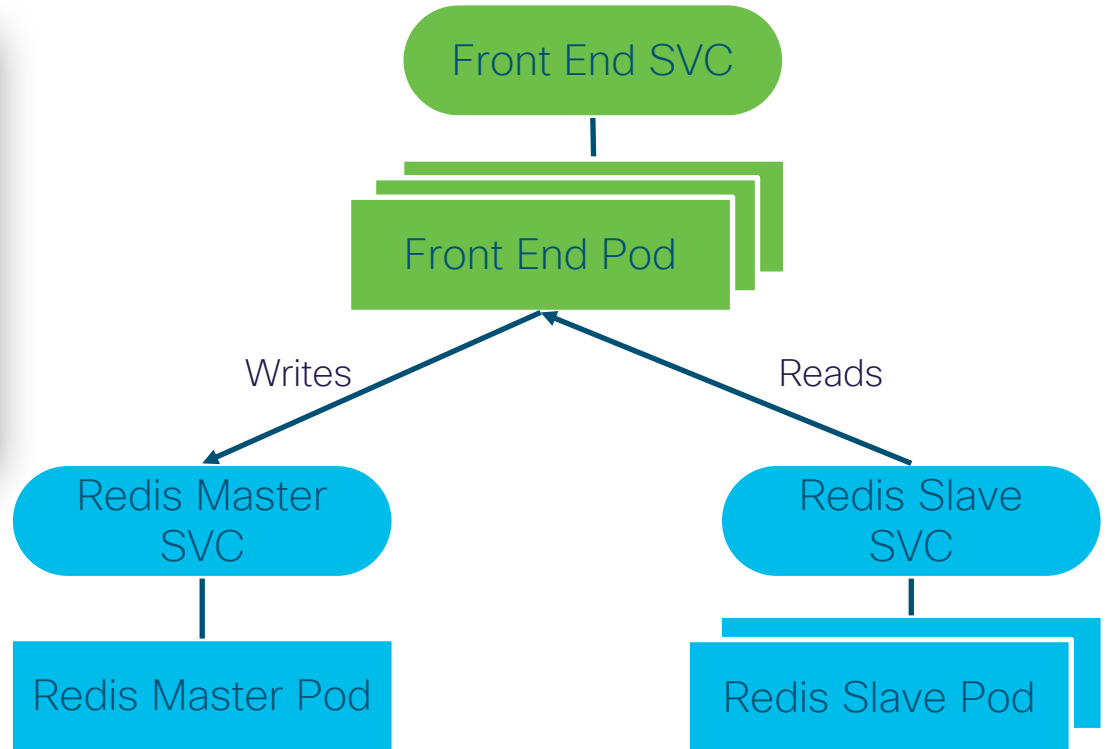
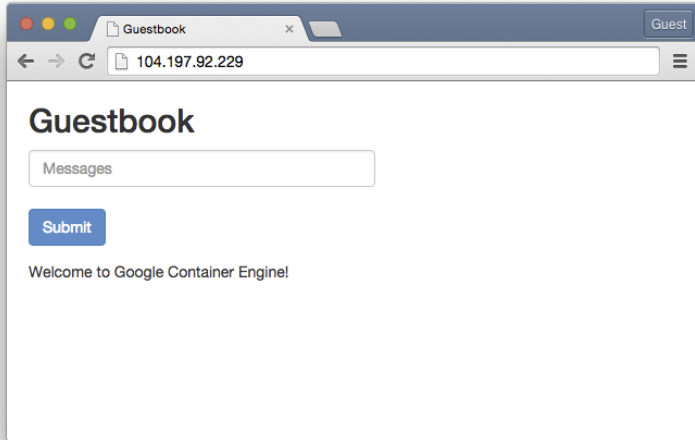
- Cisco is a major contributor to Kubeflow
 - Data pipeline operational improvements
- HX with GPUs + CCP 3.0 + Kubeflow = AWESOME

FaaS on K8s


FaaS vs Serverless

- Serverless is an application architecture approach
- FaaS is an underlying component that makes Serverless possible
- Serverless is to FaaS as Microservices are to Containers

Kubernetes Guestbook



Guestbook for FaaS on K8S with 3,000+ GitHub Stars



API Gateway + Business Logic
(Functions)



Database

Browser



http

html
&
.js



WebApp
Static Hosting
(Minio)

Example: Guestbook on Kubeless

Guestbook

Hello Python GUI

Submit

Hello World

Hello Again

Hello Python GUI

```
python -- -bash -- 86x9
PETERCJ0-M-63N9:python petercjo$ kubeless function call list
{"entries": [{"text": "Hello World", "_id": "5bd752252b215800187869ac", "updatedAt": 1540837925830}, {"text": "Hello Again", "_id": "5bd752ad2b2158001c7869ac", "updatedAt": 1540838061076}, {"text": "Hello Python GUI", "_id": "5bd753032b215800207869ac", "updatedAt": 1540838147601}]}
PETERCJ0-M-63N9:python petercjo$
```

Status	Method	File	Domain	Cause	Type	Transferred	Size	0 ms	2.56 s	6.12 s	7.68 s	10.24 s
200	GET	bootstrap.min.css	maxcdn.bootstrapcdn.com	stylesheet	css	19.69 KB	118.42 KB	→ 330 ms				
304	GET	jquery.min.js	ajax.googleapis.com	script	js	cached	0 B	→ 340 ms				
304	GET	bootstrap.min.js	maxcdn.bootstrapcdn.com	script	js	cached	0 B	→ 317 ms				
200	GET	list	10.10.20.209.xip.io	xhr	html	532 B	189 B	→ 289 ms				
200	OPTIONS	create	10.10.20.209.xip.io	xhr	html	341 B	0 B					→ 89 ms
200	POST	create	10.10.20.209.xip.io	xhr	html	433 B	91 B					→ 116 ms
200	GET	list	10.10.20.209.xip.io	xhr	html	569 B	282 B					→ 121 ms

7 requests | 118.97 KB / 439.64 KB transferred | Finish: 10.28 s | DOMContentLoaded: 324 ms | load: 401 ms

FaaS on K8S Landscape



OPENFAAS
vmware®-ish



Dockerfile	Required	Hidden	Hidden	None	Hidden	None
Image Repo	Required	Required	Required	None	None	None
Local Docker	Required	Required	Required	None	None	None
Base Image	Required	Required	Required	Required	None	None



Why FaaS on K8s Might Not Matter (Part 1 of 3)

Lambda Layers

Lambda Layers Use Cases

- Custom code, that is used by more than one function
- Libraries, modules, frameworks to simplify the implementation of your business logic

and
re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Why FaaS on K8s Might Not Matter (Part 2 of 3)

img2lambda



Clare Liguori

@clare_liguori

Following



I wrote a tool called `img2lambda` to take Docker images and convert them to AWS Lambda layers! 🥳

[github.com/awslabs/aws-la ...](https://github.com/awslabs/aws-la...)

Thread below on when `img2lambda` is useful 👉

```
> # img2lambda will extract Lambda layers from the Docker image, and publish them to Lambda
$ img2lambda -i lambda-php:latest
2019/01/22 14:42:33 Parsing the docker image docker-daemon:lambda-php:latest
2019/01/22 14:42:45 Image docker-daemon:lambda-php:latest has 5 layers
2019/01/22 14:42:46 Did not create a Lambda layer file from image layer sha256:3689366b03b394fe16903c2de2b26d0b2b651e049fce4de5eb365892b031d8d9 (no relevant files found)
2019/01/22 14:42:46 Did not create a Lambda layer file from image layer sha256:6fc15064359feb91e52dbb0621d23015c288a67ea4e623e68e744bcadaaa373f (no relevant files found)
2019/01/22 14:42:47 Created Lambda layer file output/layer-1.zip from image layer sha256:10f77e486bbae3046c6f78208635ffa5c1085c80424addebed1b803012985581
2019/01/22 14:42:47 Created Lambda layer file output/layer-2.zip from image layer sha256:153afa95b850335b89c216ca1a79ed8927fb9b105cf23d8bbfd7180e53035b4
2019/01/22 14:42:47 Did not create a Lambda layer file from image layer sha256:9316ee5ef40843ca2d6e5ef9801161fa91c8e927ce86ad9d5442de3994a80ee3 (no relevant files found)
2019/01/22 14:42:47 Created 2 Lambda layer files for image docker-daemon:lambda-php:latest
```

GIF

3:54 PM - 22 Jan 2019

https://twitter.com/clare_liguori/status/1087861037712400385

Why FaaS on K8s Might Not Matter (Part 3 of 3)

Add Outpost

If I could run Lambda on Outposts in my DC and use `img2lambda` to port container images to Lambda layers

What do I need FaaS on K8s for?

FaaS on K8s Summary

- There are 6 FaaS on K8s projects with more than 3,000 GitHub stars
 - Including entries from IBM, Google, and Oracle
- It is early for FaaS on K8s
 - Very few people using it that aren't also building it
 - Lots of tribal knowledge
- It might not matter
 - Will Outpost + Lambda crush these efforts? Too soon to say, but possible

K8s, What's Next? Takeaways

- Can AWS figure out latency issues for Outpost?
- Service Mesh and API Gateways improve K8s
- Cisco Kubeflow expertise -
> HX GPUs
- FaaS on K8s is still early and might not matter