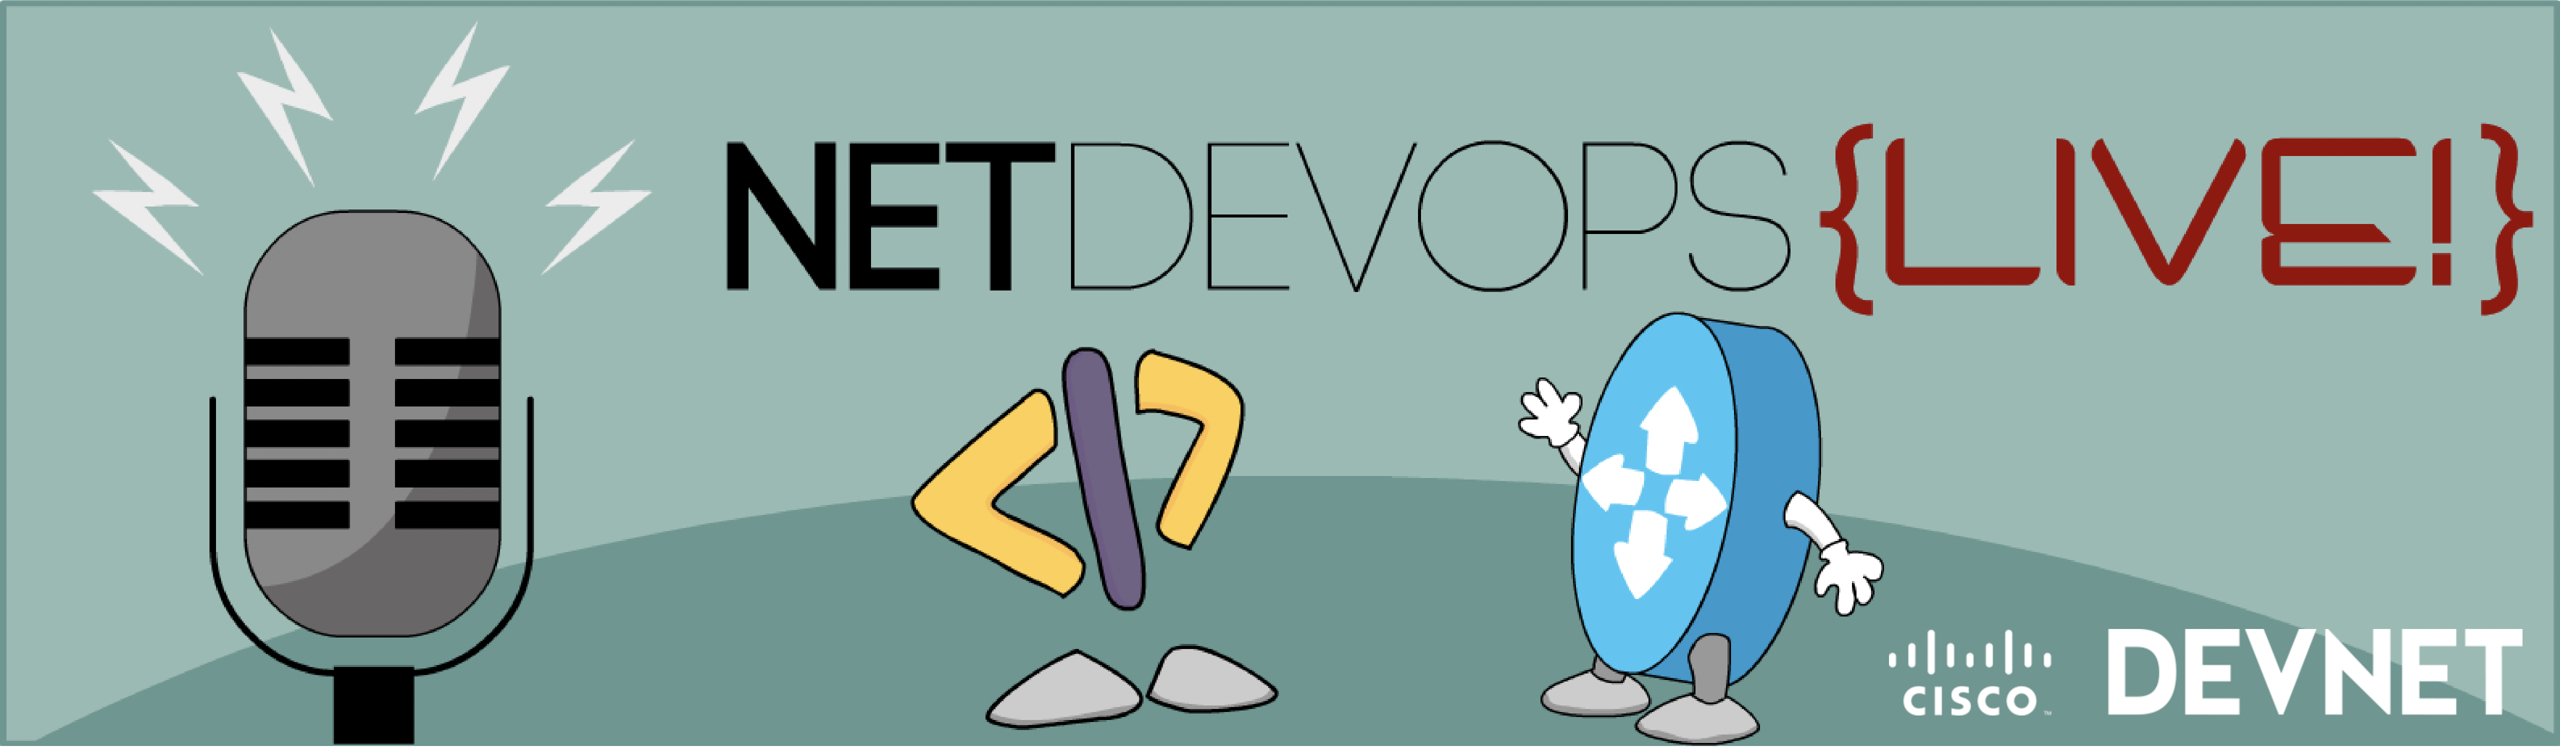

Exploring the ACI networking plugin for Kubernetes

Hank Preston, ccie 38336 R/S
Developer Advocate, DevNet
Twitter: @hfpreston

Season 1, Talk 7

https://developer.cisco.com/netdevops/live

What are we going to talk about?

- Kubernetes Basics
  - What is Kubernetes?
  - Key objects in Kubernetes
  - Networking in Kubernetes
- ACI + Kubernetes
  - What do you get?
  - A bit on how it works
- ACI + Kubernetes Demonstration

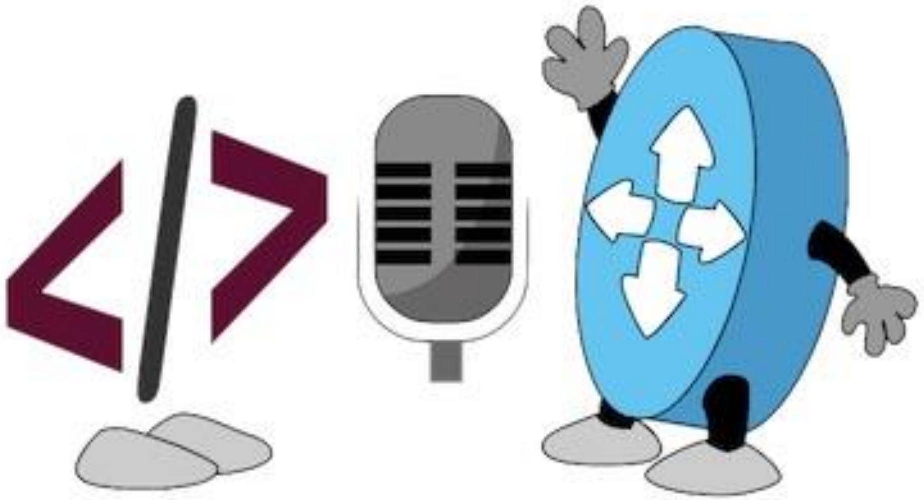

Kubernetes Basics

Container Orchestration 101

- Bring multiple hosts together and make them part of a cluster
- Schedule containers to run on different hosts
- Help containers running on one host reach out to containers running on other hosts in the cluster
- Bind containers and storage
- Bind containers of similar type to a higher-level construct, like services, so we don't have to deal with individual containers
- Keep resource usage in-check, and optimize it when necessary
- Allow secure access to applications running inside containers.

Kubernetes

- Kubernetes is an open source Container Orchestration system for automating deployment, scaling and management of containerized applications.
- It was inspired by the Google Borg System and with its v1.0 release in July 2015, Google donated it to the Cloud Native Computing Foundation (CNCF).
- Generally, Kubernetes has new releases every three months.

Kubernetes & Docker

- Kubernetes uses Docker to execute/run the containers
- Kubernetes adds, on top of Docker, all the intelligence and features of an orchestrator

Kubernetes Architecture

- At a very high level, Kubernetes has the following main components:
  - One or more Master Nodes
  - One or more Worker Nodes
  - Distributed key-value store, like etcd.

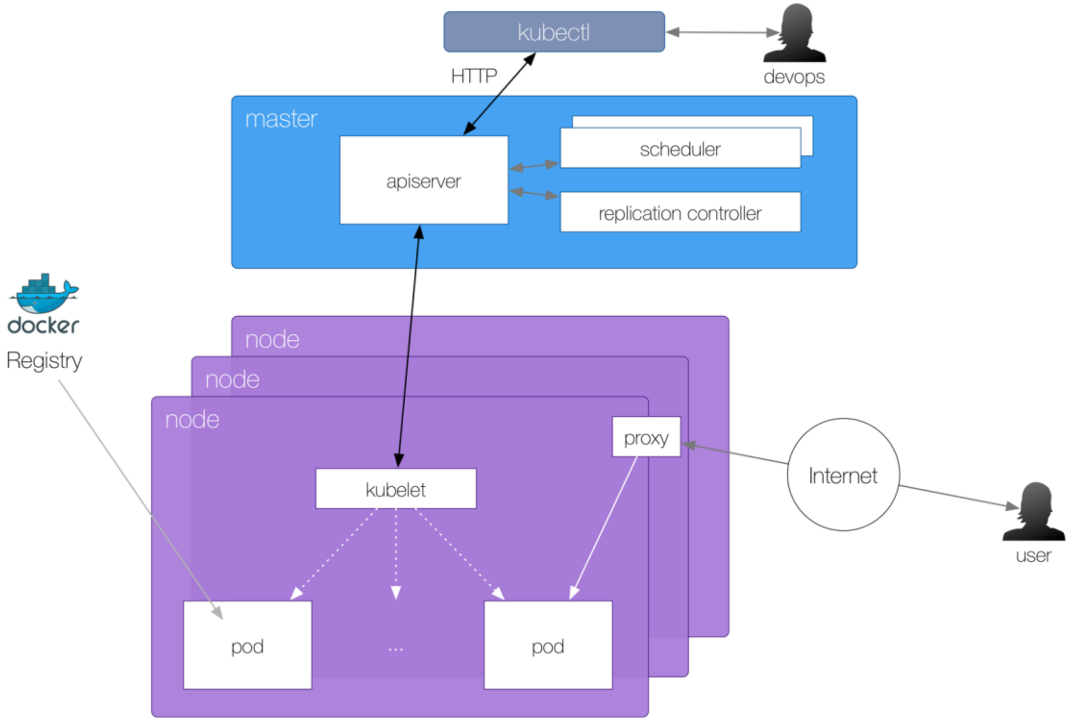

Kubernetes Components – Master Node

- The Master Node is responsible for managing the Kubernetes cluster.
- Master node access methods are CLI, GUI or APIs.
- For fault tolerance, there can be more than one Master Node.

Master Node

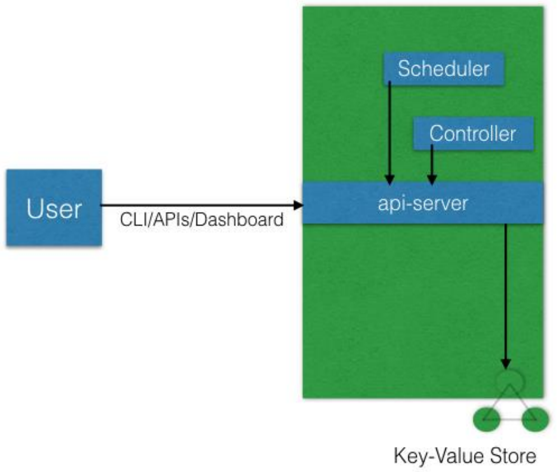

Kubernetes Components – Worker Node

- A Worker Node is a machine (VM, physical server, etc.)
- Runs the containers using “pods”
- Controlled by the Master Node.

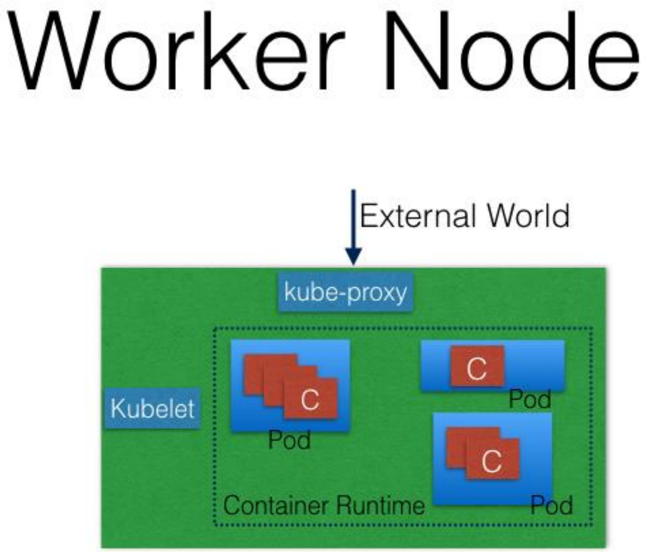

Kubernetes Key Objects

Conceptual and just enough for this lab…

- A Deployment represents a Micro Service description
- A Pod is an instantiation of the Deployment (typically a “containers”)
- A Service provides a single entry-point to a Deployment (think load balancer)
  - Cluster IPs are for intra-Kubernetes connections
  - External IPs are for extra-Kubernetes connections

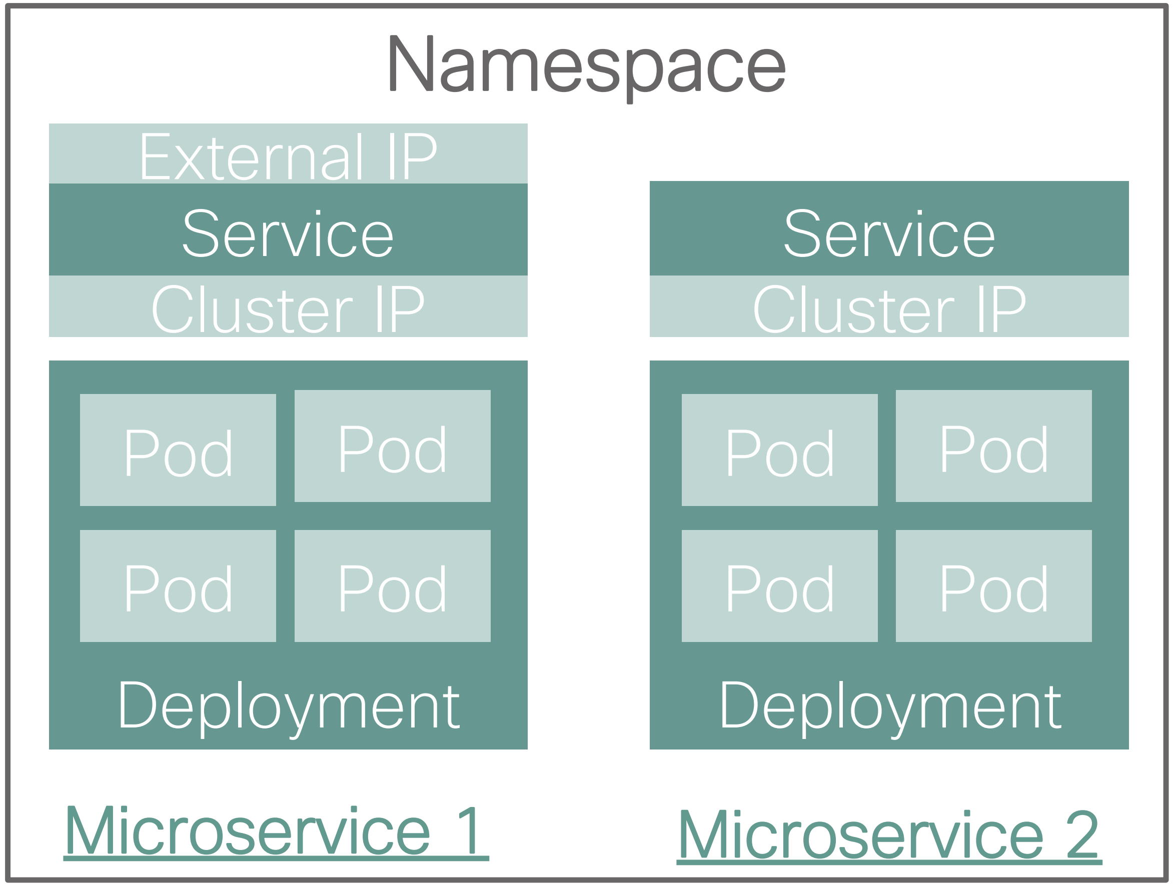

- Namespace is an organizational construct

Kubernetes Annotations

- Meta-data attached to Kubernetes Objects
- Can be attached to ANY object
- Not directly used by Kubernetes, available for plugins and other tooling

```
myhero
  Name:           myhero
  Labels:         <none>
  Annotations:
    opflex.cisco.com/endpoint-group
      ={"tenant": "kubesbx04",
        "app-profile": "kubernetes",
        "name": "ns-myhero"}

  Status:         Active

  No resource quota.

  No resource limits.
```

A tale of two standards…

- Custom network driver (CNM)
  - Proposed by Docker
  - Plugin-based
  - Supports Only Docker
  - Containers con join 1 or more networks
  - Supports namespace isolation
  - Integrates with IPAM
  - Complex

- Container network interface (CNI)
  - Proposed by CoreOS
  - Plugin-based
  - Multiple runtime (Docker, LXC etc..)
  - Containers con join 1 or more networks
  - Supports namespace isolation
  - Integrates with IPAM
  - Simple

Kubernetes choose… CNI

http://blog.kubernetes.io/2016/01/why-Kubernetes-doesnt-use-libnetwork.html

ACI + Kubernetes

Cisco ACI and Container Integration

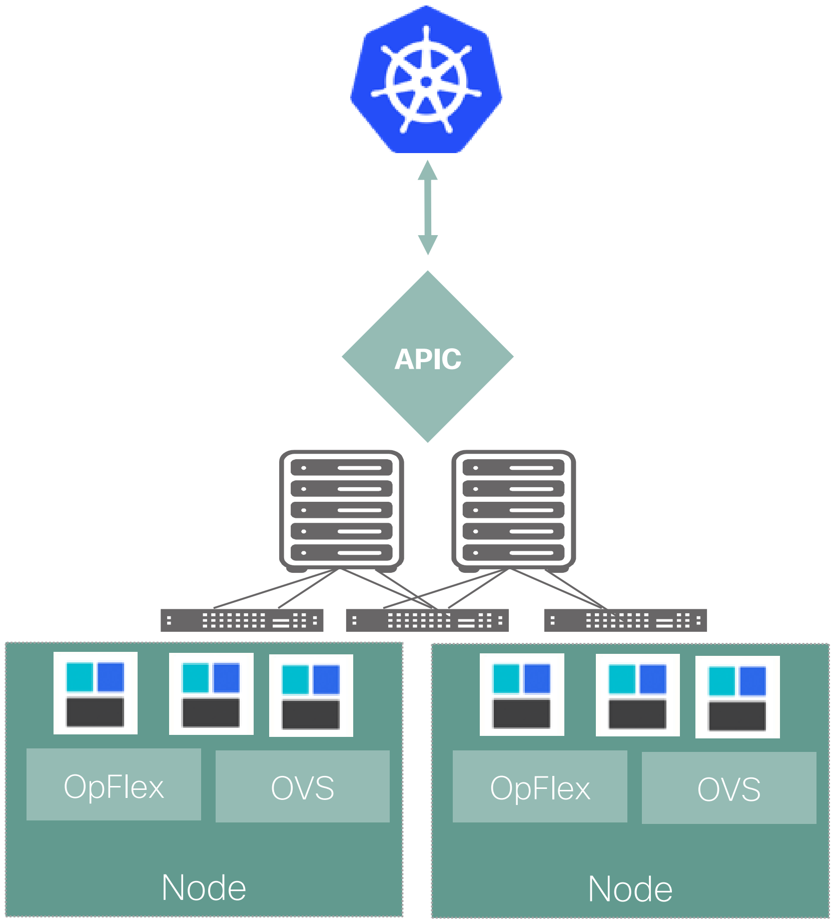

ACI and Containers

- Unified networking: Containers, VMs, and bare-metal
- Micro-services load balancing integrated in fabric for HA / performance
- Secure multi-tenancy and seamless integration of Kubernetes network policies and ACI policies
- Visibility: Live statistics in APIC per container and health metrics

ACI Network Plugin for Kubernetes

Native Security Policy Support

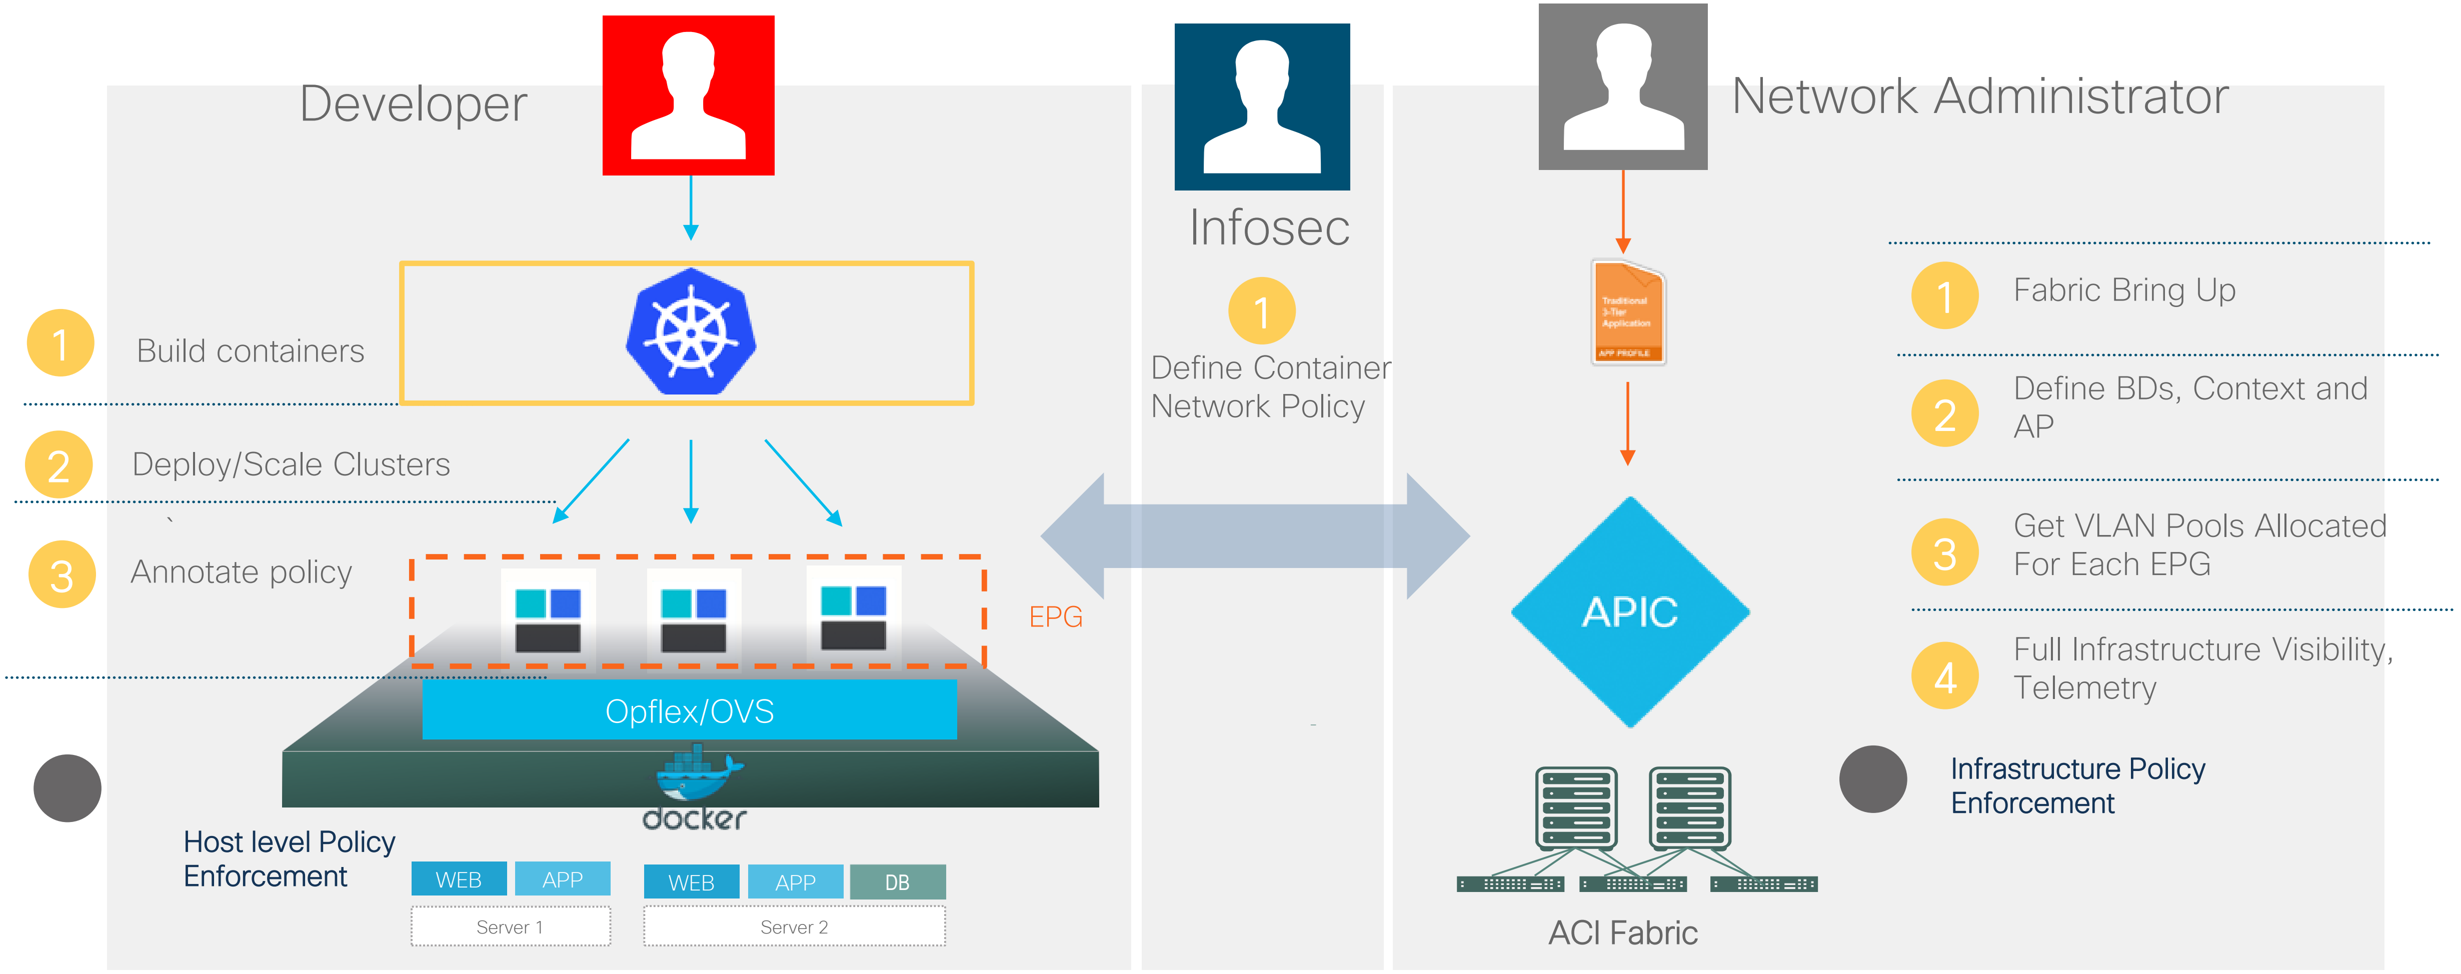

Developer

1 Build containers

2 Deploy/Scale Clusters

3 Annotate policy

EPG

Opflex/OVS

Host level Policy Enforcement

WEB	APP
Server 1	

WEB	APP	DB
Server 2		

Infosec

1 Define Container Network Policy

Network Administrator

1 Fabric Bring Up

2 Define BDs, Context and AP

3 Get VLAN Pools Allocated For Each EPG

4 Full Infrastructure Visibility, Telemetry

Infrastructure Policy Enforcement

APIC

ACI Fabric

ACI VMM Domain for Kubernetes

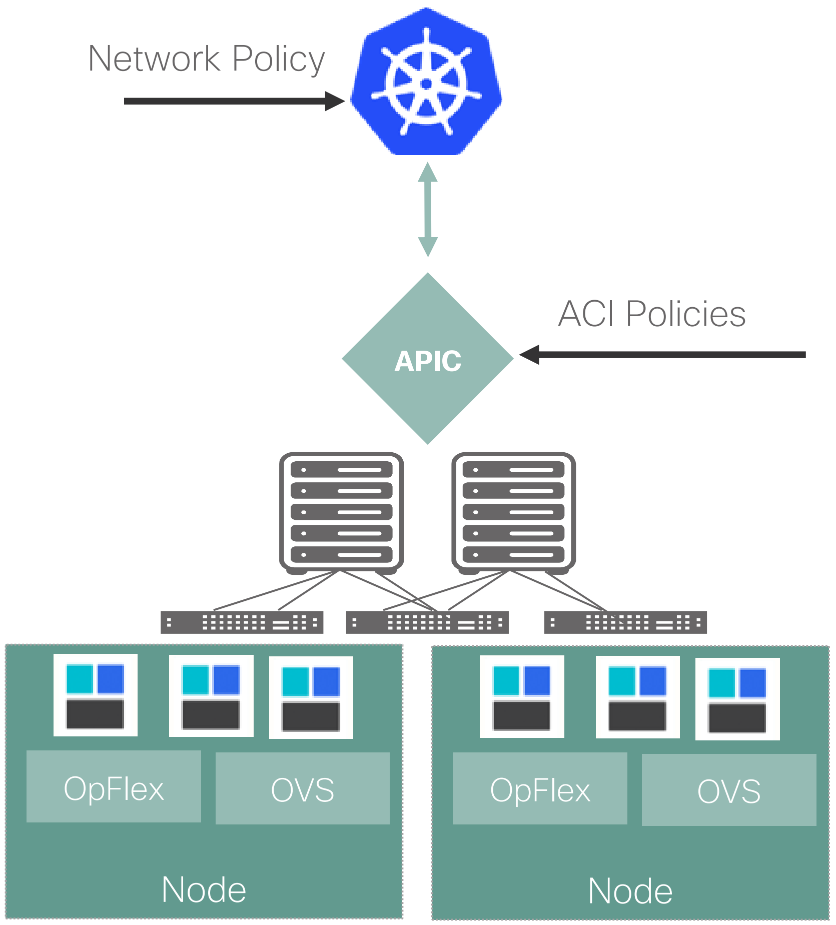

Technical Description

- Network policies of Kubernetes supported using standard upstream format but enforced through OpFlex / OVS using APIC Host Protection Profiles
- Kubernetes app configurations can be moved without modification to/from ACI and non-ACI environments
- Embedded fabric and virtual switch load balancing
  - PBR in fabric for external service load balancing
  - OVS used for internal service load balancing
- VMM Domain for Kubernetes
  - Stats per namespace, deployment, service, pod
  - Physical to container correlation

ACI CNI Plugin Components

- aci-containers-controller
  - Monitors Kubernetes application state & ACI configuration and ensures they are in sync.
- aci-containers-host
  - Manage node level configurations on each Kubernetes node.
- aci-containers-openvswitch
  - Provides the actual networking functions on each node.

```
NAME                                  READY     STATUS    RESTARTS   AGE       IP            NODE
aci-containers-controller-3029831268-nj1hq   1/1       Running   0          20h       172.20.0.39   sbx38kube03.localdomain
aci-containers-host-9s4tm             3/3       Running   0          20h       172.20.0.39   sbx38kube03.localdomain
aci-containers-host-bp01p             3/3       Running   0          20h       172.20.0.38   sbx38kube02.localdomain
aci-containers-host-gzvdj             3/3       Running   0          20h       172.20.0.37   sbx38kube01.localdomain
aci-containers-openvswitch-0klgg      1/1       Running   0          20h       172.20.0.37   sbx38kube01.localdomain
aci-containers-openvswitch-ds64p      1/1       Running   0          20h       172.20.0.39   sbx38kube03.localdomain
aci-containers-openvswitch-vcr7c      1/1       Running   0          20h       172.20.0.38   sbx38kube02.localdomain
```

Mapping Network Policy and EPGs

Cluster Isolation

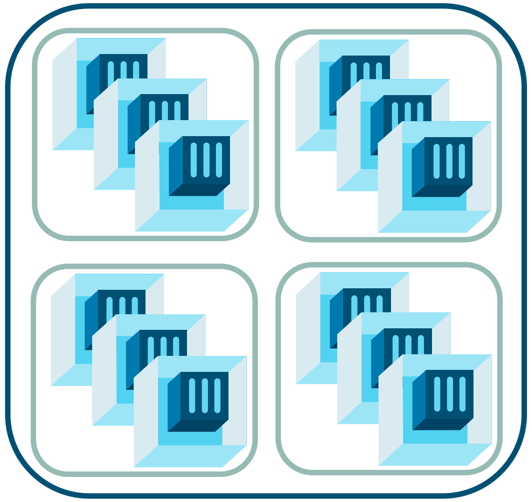

Single EPG for entire cluster.

No need for any internal contracts.
(Default behavior)

Namespace Isolation

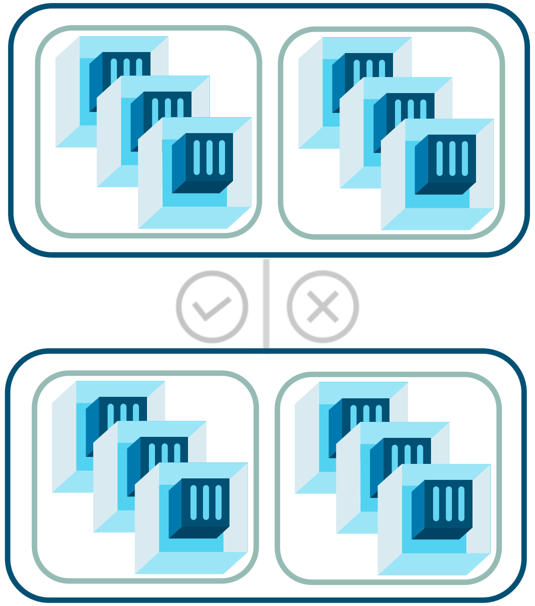

Each namespace is mapped to its own EPG.

Contracts for inter-namespace traffic.

Deployment Isolation

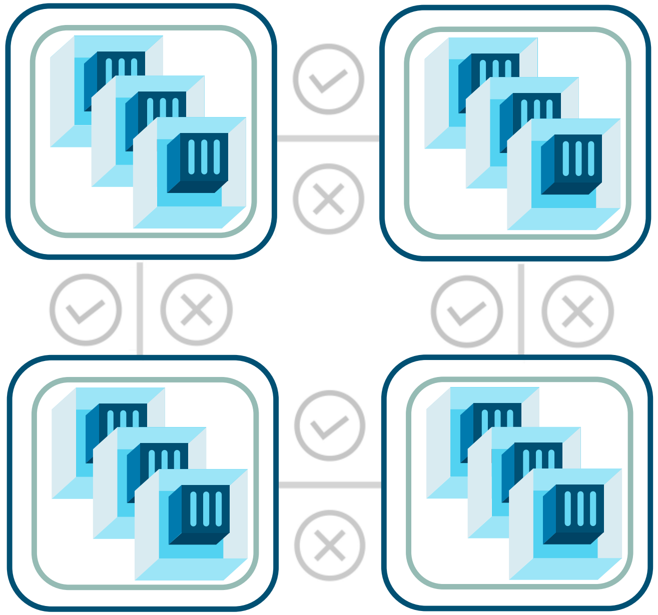

Each deployment mapped to an EPG

Contracts tightly control service traffic

Key Map | EPG | NetworkPolicy | ✓ | ⊗ Contract

ACI + Kubernetes Demonstration

Summing up

What did we talk about?

- Kubernetes Basics
  - What is Kubernetes?
  - Key objects in Kubernetes
  - Networking in Kubernetes
- ACI + Kubernetes
  - What do you get?
  - A bit on how it works
- ACI + Kubernetes Demonstration

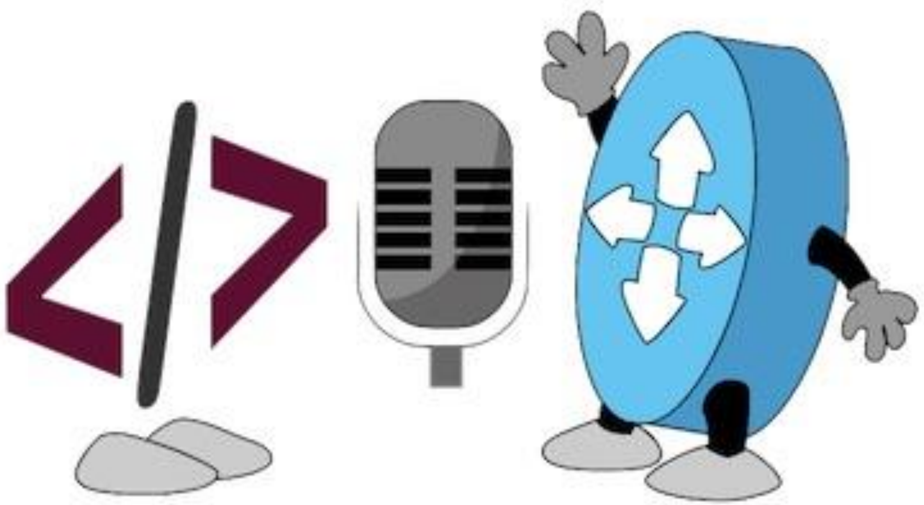

Webinar Resource List

- Docs and Links
  - Deploying Kubernetes in the Enterprise with Cisco ACI - BRKACI-2505
  - Cisco ACI and Kubernetes Integration Guide
  - Cisco ACI and OpenShift Integration Guide
- Learning Labs
  - Exploring the ACI CNI plug-in for Kubernetes http://cs.co/lab-acik8s
  - DevOps 101 http://cs.co/lab-devops-apps
- DevNet Sandboxes
  - ACI and Kubernetes Sandbox http://cs.co/sbx-acik8s
- Code Samples
  - http://cs.co/code-acik8s

NetDevOps Live! Code Exchange Challenge

developer.cisco.com/codeexchange

Deploy a sample application to Kubernetes/ACI with Deployment Isolation. Provide application definition for Kubernetes and ACI.

Example: Find sample applications at https://github.com/kubernetes/examples

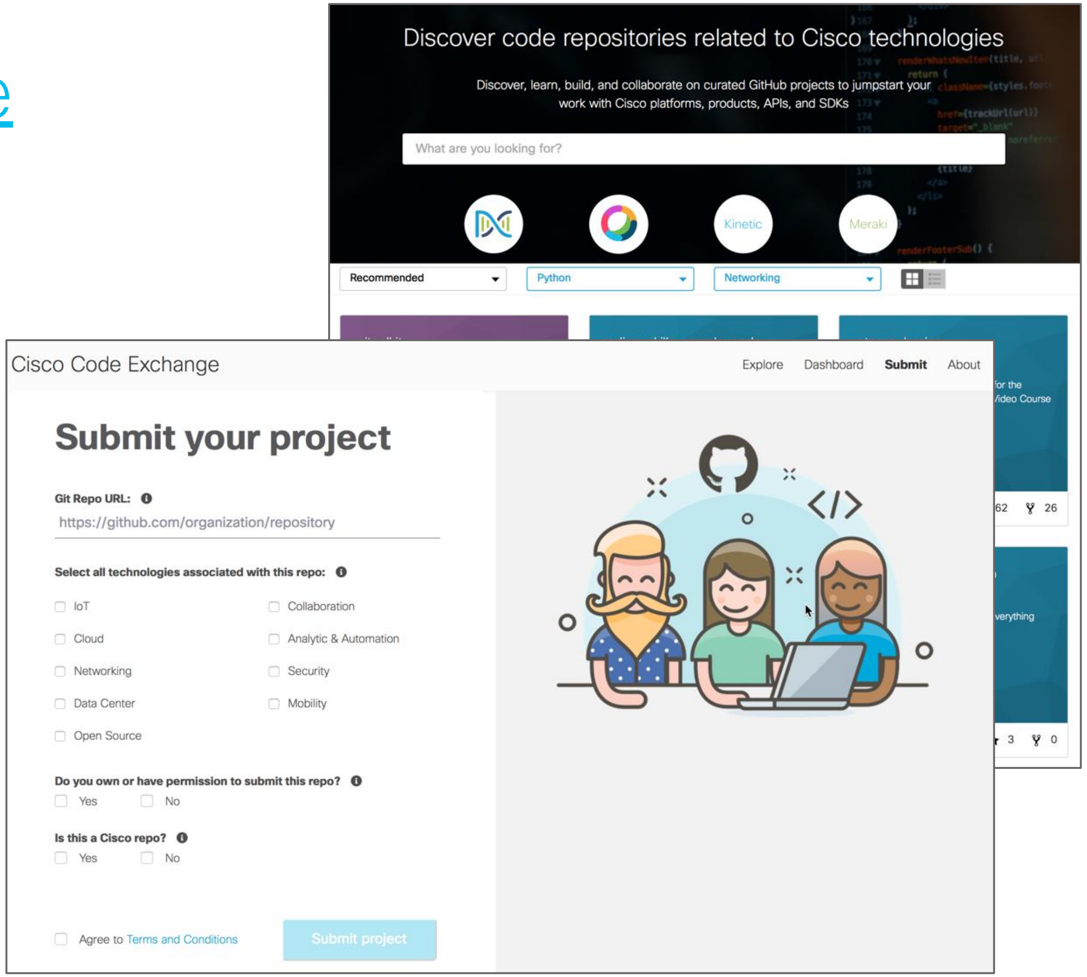

Discover code repositories related to Cisco technologies

Discover, learn, build, and collaborate on curated GitHub projects to jumpstart your work with Cisco platforms, products, APIs, and SDKs

What are you looking for?

Cisco Code Exchange

Submit your project

Git Repo URL: https://github.com/organization/repository

Select all technologies associated with this repo:

□ IoT
□ Cloud
□ Networking
□ Data Center
□ Open Source
□ Collaboration
□ Analytic & Automation
□ Security
□ Mobility

Do you own or have permission to submit this repo?
□ Yes □ No

Is this a Cisco repo?
□ Yes □ No

□ Agree to Terms and Conditions

Looking for more about NetDevOps?

- NetDevOps on DevNet
  developer.cisco.com/netdevops
- NetDevOps Live!
  developer.cisco.com/netdevops/live
- NetDevOps Blogs
  blogs.cisco.com/tag/netdevops
- Network Programmability Basics Video Course
  developer.cisco.com/video/net-prog-basics/

Got more questions? Stay in touch!

Hank Preston

hapresto@cisco.com

@hfpreston

http://github.com/hpreston

CISCO

DEVNET

LEARN CODE INSPIRE CONNECT

developer.cisco.com

@CiscoDevNet

facebook.com/ciscodevnet/

http://github.com/CiscoDevNet

NETDEVOPS{LIVE!}

cisco DEVNET

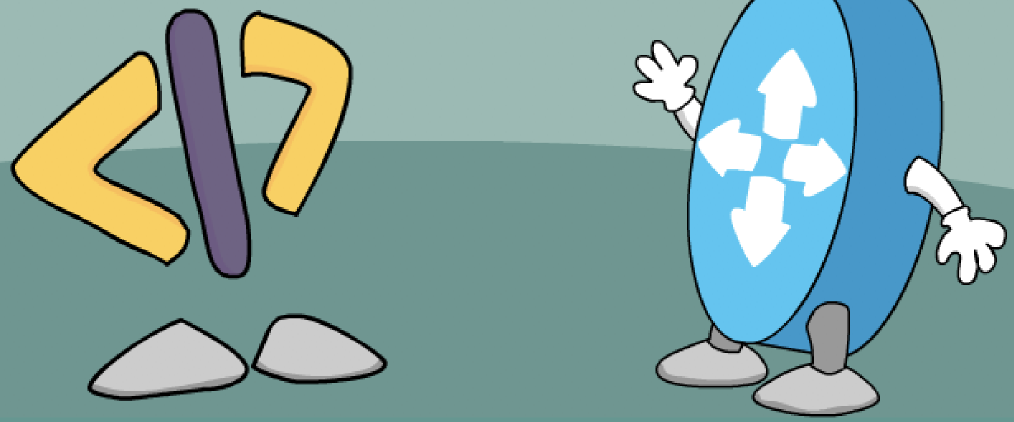

https://developer.cisco.com/netdevops/live

@netdevopslive