



CLOUD NATIVE
COMPUTING FOUNDATION

IoT Edge Working
Group

Presenter



Dejan Bosanac
Red Hat
@dejanb

IoT Edge Working Group

- Community to discuss all the topics around IoT and Edge Computing with K8S
- What kind of use cases, workloads and applications are common to running K8s on the Edge?
- Why and how Kubernetes can be a good fit to these Edge scenarios?
- What are the key challenges for deploying of Kubernetes to Edge today?

Agenda

- Brief introduction to Edge Computing and K8s
- Edge workloads and how to run them
- Workload challenges
- Ecosystem
- Future

What is Edge Computing?

- Edge - everything that is not cloud
- Effort to deploy workloads closer to the users and devices

Why Edge Computing?

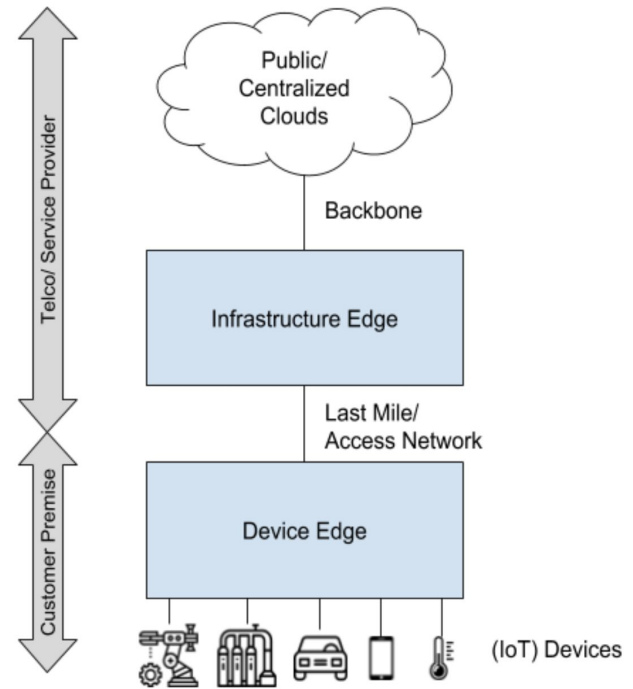
- Latency
- Availability
- Data locality
- ...

Use Cases

- Telco Infrastructure
- Large scale IoT and IIoT
- Video games
- VR/AR
- AI/ML
-

Edge types

- Infrastructure Edge
 - Deploy whole clusters on Edge sites
 - Hybrid-clouds
 - Federated clusters
- Device Edge
 - Deploy cluster nodes outside of the cloud

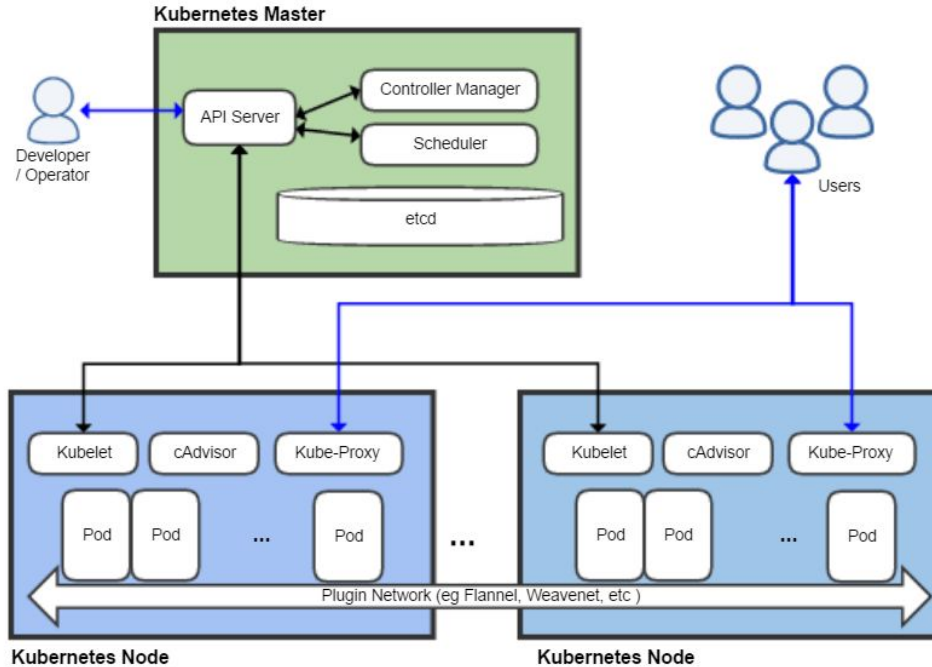


Source: Icons from <https://www.flaticon.com/free-icon>

Why Kubernetes?

- Kubernetes API
 - Developers mindshare and tooling
 - Same workloads can move between Cloud and Edge
- Kubernetes Architecture
 - No need to reinvent it

Kubernetes Architecture



Edge Challenges

- Infrastructure
 - How to manage resources (nodes and clusters) on the Edge?
- Control plane
 - How to manage workloads on the Edge?
- Data plane
 - How Edge sites communicates with the cloud and between themselves?



Edge Workloads

Kubernetes for the Edge



KubeCon

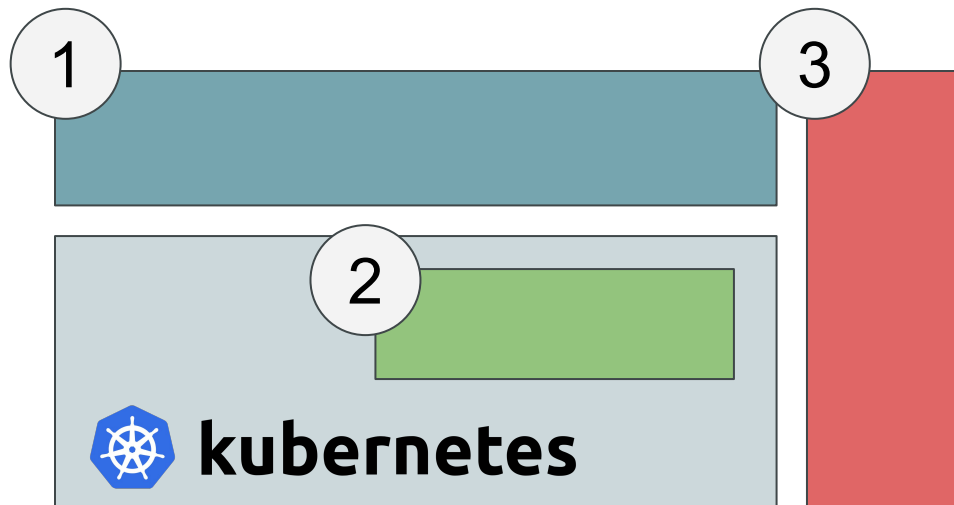


CloudNativeCon

North America 2018

How does Kubernetes interact with the Edge?

1. Edge workloads that run ON Kubernetes
2. Edge challenges mitigated BY Kubernetes
3. Edge capabilities not easily serviceable by Kubernetes



Kubernetes features

Edge workload

```
tolerations:  
- key: "edge"  
  operator: "Equal"  
  value: "true"  
  effect: "NoSchedule"
```

Cloud workload

```
tolerations:  
- key: "cloud"  
  operator: "Equal"  
  value: "true"  
  effect: "NoSchedule"
```

EdgeNode

```
kubectl taint nodes EdgeNode edge=true:NoSchedule
```

Edge workloads - Why?

- Data ingestion and processing
 - Protocol conversion
 - Data preprocessing
- Reliability and availability
 - Buffer and batch
 - Caching
- Latency
 - Edge functions
 - Compute offloading
 - Machine learning

Protocol conversion

- Network level
 - Converting non-IP protocols to TCP/IP based ones
 - Modbus in IIoT
 - Bluetooth in consumer IoT
 - Usually converting to some widely used messaging protocol
 - MQTT
 - AMQP
 - HTTP



- Kubernetes supports "device plugins"
- Taints and tolerances can be used for scheduling to appropriate nodes
- New concepts for easier access to interfaces
 - <https://www.networkservicemesh.io/>

Data preprocessing

- Convert data to general structured messages
- Normalize data structure
 - Vorto, LWM2M
- Data analytics
 - Send only relevant data
 - Combine multiple sources
- Add metadata
 - Location
 - Identity
 - Security



Generic Kubernetes workloads
Needs to be properly containerized and
orchestrated on the Edge nodes

Reliability and high availability

- Buffer and batch
 - Store and forward
 - Brokers on Edge nodes
- Caching
 - Local databases on Edge nodes
 - Sync data with the cloud and other Edge nodes



Edge Nodes/Clusters may have limited storage volumes to hold data until it can be uploaded

Latency: Functions

- React locally on sensor or scheduled events



- Possible CNCF projects collaboration
 - Cloud Events - <https://cloudevents.io/>
 - Knative - <https://github.com/knative/>

Latency

- Compute offload
 - Schedule resource intensive tasks on the dedicated hardware on the Edge
 - Example AR/VR renderings
- Machine learning
 - Cloud trained models - executed on the Edge
 - Edge specific training (environment and data policies)



Taints and tolerances can be used for scheduling to appropriate nodes (e.g. GPU availability)



CLOUD NATIVE
COMPUTING FOUNDATION

Device Edge Challenges

Workload challenges


- **Problem:** Limited node resources
- **Solution:** Workload prioritization

- **Problem:** Unsecure, unreliable, limited network
- **Solution:** Network policies and traffic shaping


Workload prioritization - Why

- Limited number of nodes on the Edge
- No autoscaling
- Workloads with wide range of priorities
- Adds more emphasis on prioritization

Cloud



Edge



Kubernetes prioritization toolkit

Prioritization

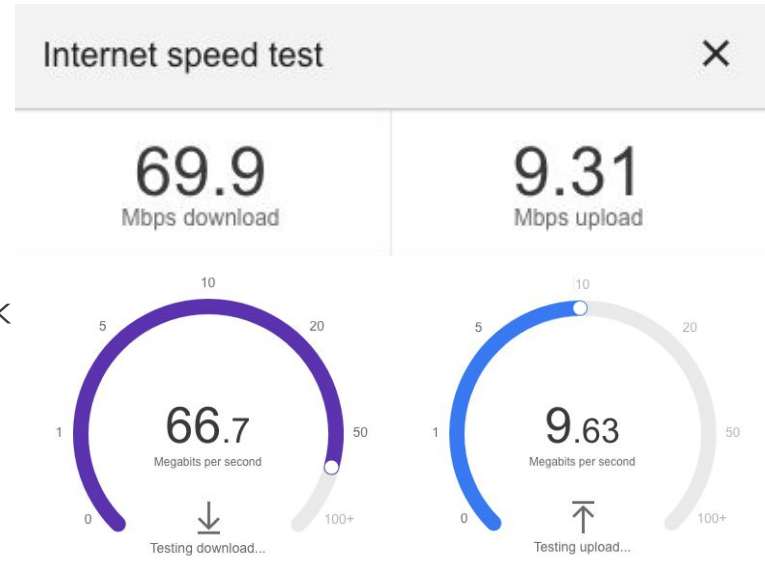
- Ranking of priority classes
- Input to pre-emption logic
- Applied to a pod, but acted on by node
- Different from resource based eviction

Quality of Service

- Three levels
 - Guaranteed
 - Burstable
 - Best Effort
- These are implicit from pod spec
- Is NOT considered for preemption
- IS considered in the case of eviction
- preemption != eviction

Traffic shaping - Why

- Managing bandwidth
- Network capacity can be limited
- Different workloads should have different network
- Related to "Workload Prioritization"



Traffic shaping

Policy

- Deals with what traffic is allowed
- Applied via Network Plugin
- Creates NetworkPolicy resource
- Based on 'cluster-external' IPs
- Based on SRC/DST and port
 - src/dst can be specified several ways
 - May be subject to cluster environment

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: test-network-policy
  namespace: default
spec:
  podSelector:
    matchLabels:
      role: db
  policyTypes:
  - Ingress
  - Egress
  ingress:
  - from:
    - ipBlock:
        cidr: 172.17.0.0/16
        except:
        - 172.17.1.0/24
    - namespaceSelector:
        matchLabels:
          project: myproject
    - podSelector:
        matchLabels:
          role: frontend
  ports:
  - protocol: TCP
    port: 6379
  egress:
  - to:
    - ipBlock:
        cidr: 10.0.0.0/24
  ports:
```

Traffic shaping

Bandwidth

- Effected by a set of layers
- Does not manage bandwidth cluster wide

Pod: bandwidth annotations

CNI: Bandwidth Plugin

`tc` (Traffic Control)

Linux: Network Namespace

Traffic shaping

Bandwidth: CNI

- Can be enabled as a plugin without specific limits

```
{
  "type": "bandwidth",
  "capabilities": {"bandwidth": true}
}
```

- Can be chained to a specific network interface and limit interface bandwidth use

```
{
  "cniVersion": "0.3.1",
  "name": "mynet",
  "plugins": [
    {
      "type": "ptp",
      "ipMasq": true,
      "mtu": 512,
      "ipam": {
        "type": "host-local",
        "subnet": "10.0.0.0/24"
      },
      "dns": {
        "nameservers": [ "10.1.0.1" ]
      }
    },
    {
      "name": "slowdown",
      "type": "bandwidth",
      "ingressRate": 123,
      "ingressBurst": 456,
      "egressRate": 123,
      "egressBurst": 456
    }
  ]
}
```

Traffic shaping



KubeCon



CloudNativeCon

North America 2018

Bandwidth: Pod Spec

```
{
  "kind": "Pod",
  "metadata": {
    "name": "iperf-slow",
    "annotations": {
      "kubernetes.io/ingress-bandwidth": "10M",
      "kubernetes.io/egress-bandwidth": "10M"
    }
  }
}
```





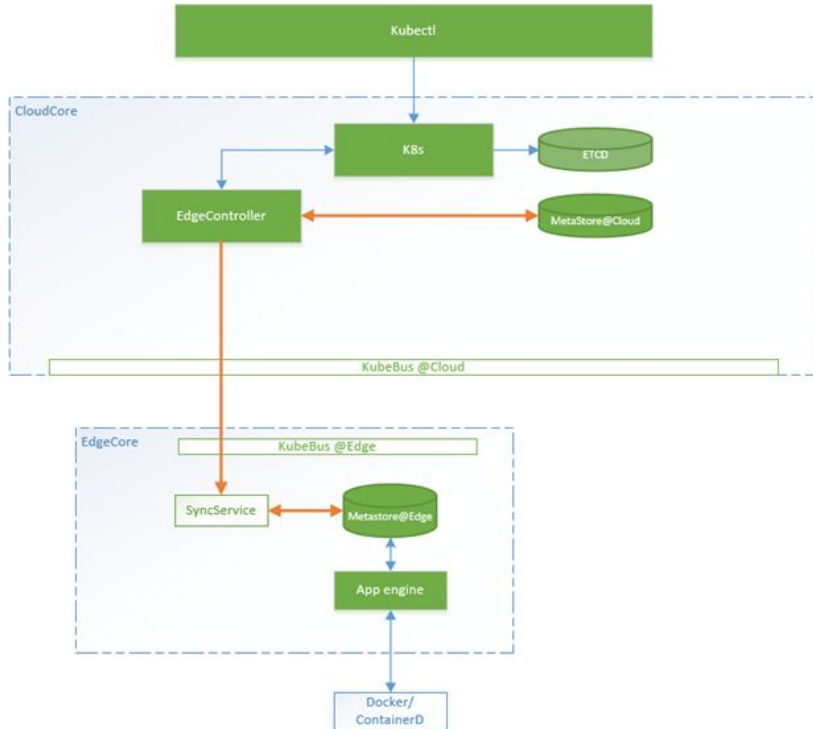
CLOUD NATIVE
COMPUTING FOUNDATION

Beyond Kubernetes

Missing solutions for ...

- Hard Edge cluster deployment
- Unreliable network between control plane and nodes
- Nodes with constrained resources

KubeEdge



- <https://kubedge.io/en/>
- IoT Edge platform built on Kubernetes
- Constrained resources with lightweight agent
- Edge autonomy with KubeBus

Virtual Kubelet

- <https://virtual-kubelet.io/>
- Framework for implementing alternative Kubelet implementations
- Problem: Deploying nodes on constrained devices
- Example: Azure IoT Edge Connector for Kubernetes

k3s

- <https://k3s.io>
- Lightweight Kubernetes distribution suitable for the Edge
- Problem: Infrastructure plane ... deploying clusters to the Edge

Eclipse ioFog

- <https://iofog.org/>
- Microservices framework for the Edge
- A lot of helpful tools and concepts to deploy containerized workloads to the Edge
- Moving towards Kubernetes

Future activities

- Dig deeper into specific topics
 - Security
 - Messaging
 - ...

Questions?

Regular Work Group Meeting: Wednesdays at 10:00am/pm Pacific (bi-weekly)

- [Meeting notes and agenda](#)

Link to join the group

- <https://groups.google.com/forum/#!forum/kubernetes-wg-iot-edge>

Link to join Slack

- <https://kubernetes.slack.com/messages/wg-iot-edge>

White Paper

- <http://bit.ly/iot-edge-whitepaper>
 - Workloads being considered
 - Technical challenges
 - Available architectural solutions

KubeCon EU 2019

- <https://kccnceu19.sched.com/event/MPI4/intro-deep-dive-kubernetes-iot-edge-wg-steven-wong-vmware-cindy-xing-huawei-dejan-bosanac-red-hat>