

Navigating the Networking Maze of Kubernetes

A Journey of Discovery, Confusion, and (Hopefully) Enlightenment


aojea@google.com




FOSDEM 2024

Google

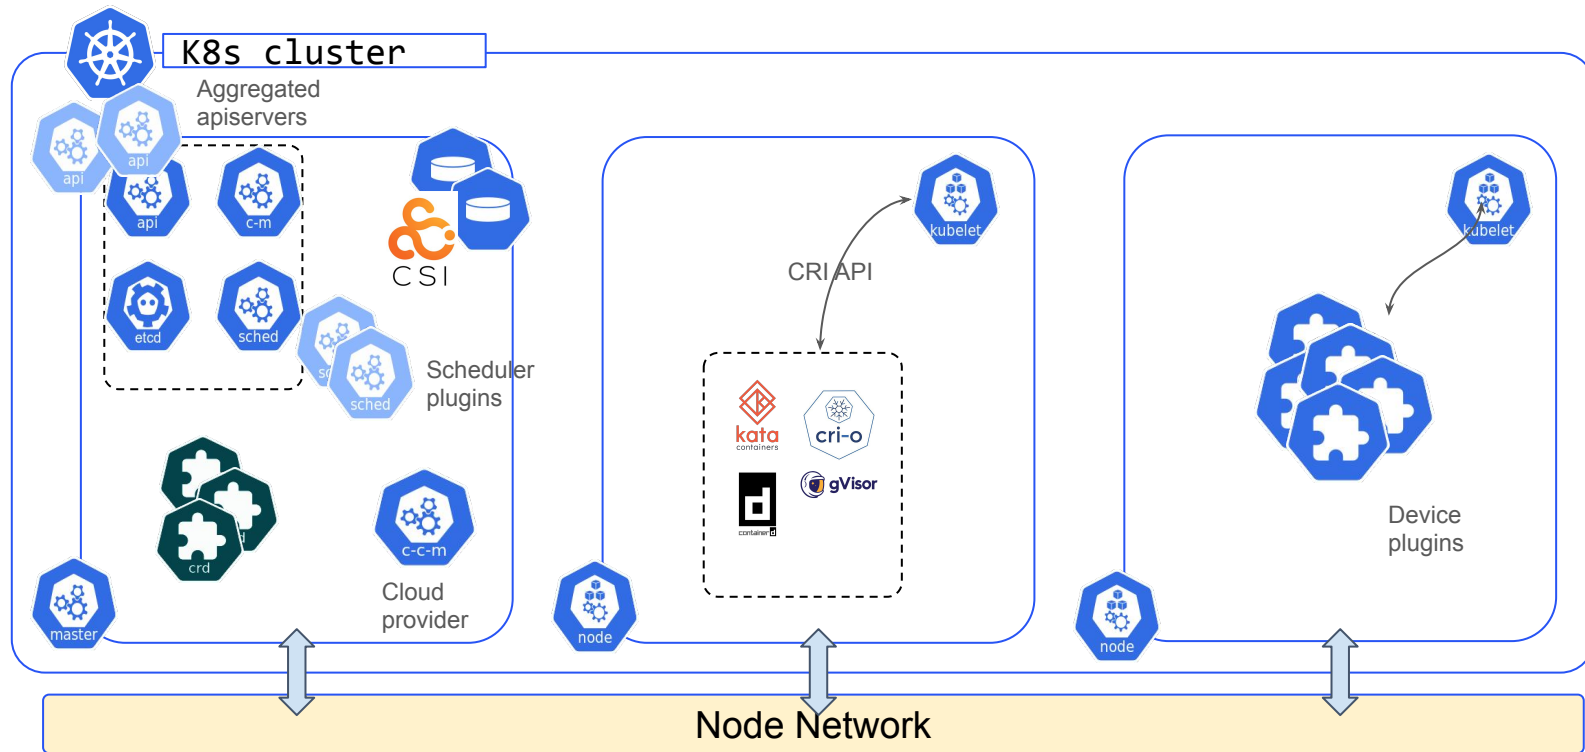
How it started



How it is going



Kubernetes Architecture: extensible and pluggable



Kubernetes Architecture: API + Semantics

Kubernetes API: IPv4

```
apiVersion: v1
kind: Node
metadata:
  name: node1
spec:
  podCIDR: 10.0.0.0/24
...
```

```
status:
  addresses:
  - address: 192.168.0.5
    type: InternalIP
```

```
apiVersion: v1
kind: Pod
metadata:
  name: app-http
  namespace: default
spec:
  containers:
  - name: app
    image: myapp:0.1
...
```

```
status:
  phase: Running
  podIP: 10.0.0.5
```

```
apiVersion: v1
kind: Service
metadata:
  name: lb-app
  namespace: default
spec:
  clusterIP: 10.96.0.13
...
```

```
type: LoadBalancer
status:
  loadBalancer:
  ingress:
  - ip: 196.23.45.23
```

community / contributors / devel / sig-testing / e2e-tests.md

lunarwhite Fix broken links in contributors guide ✓

Preview Code Blame 786 lines (684 loc) · 35 KB

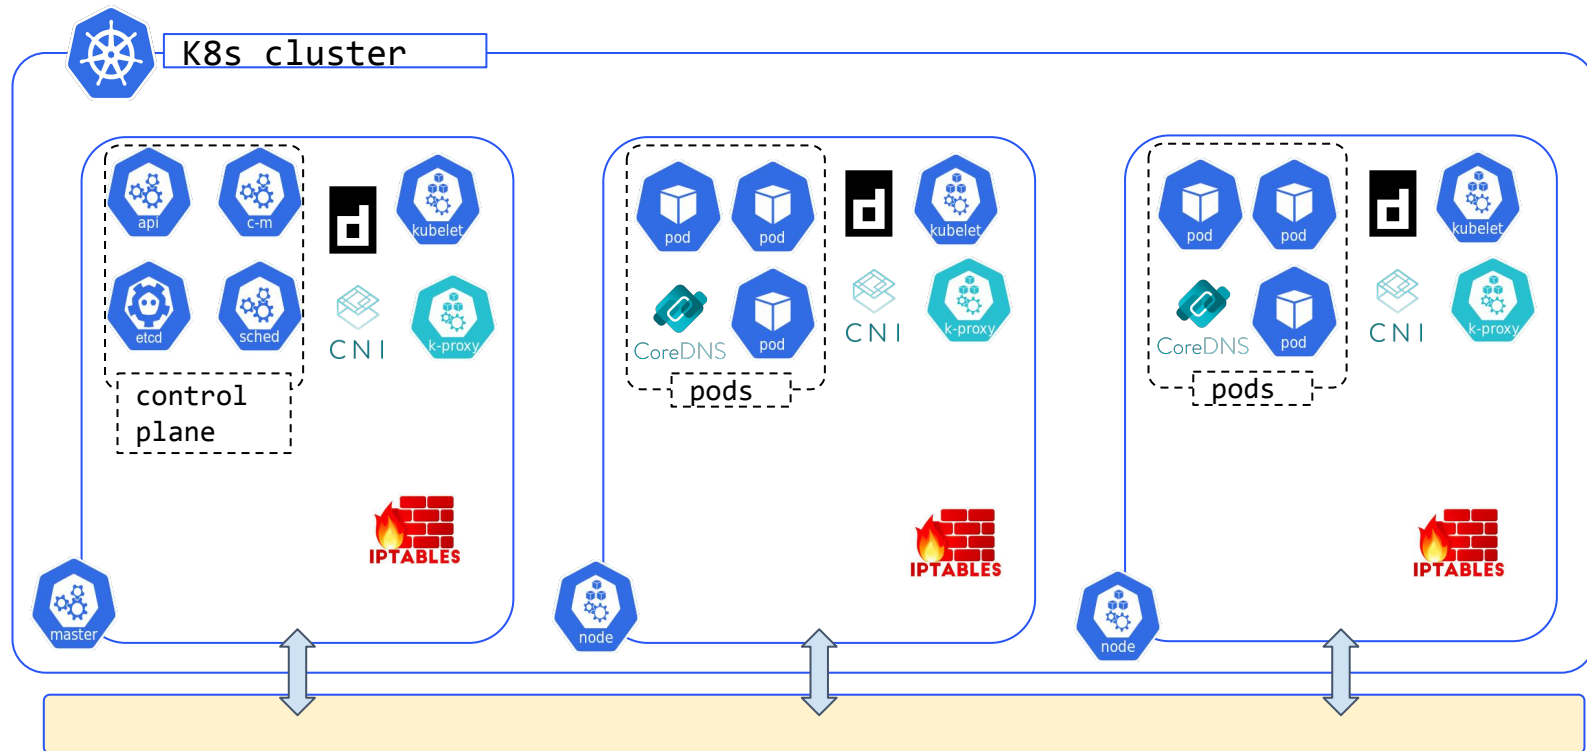
End-to-End Testing in Kubernetes

Table of Contents

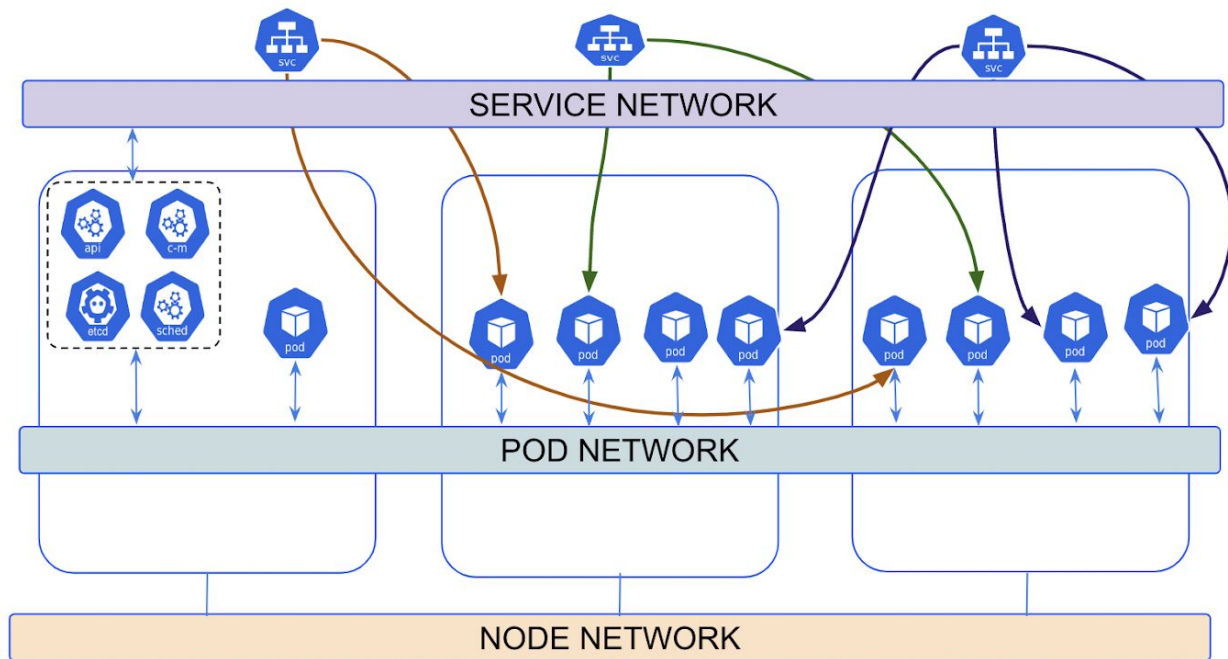
- [End-to-End Testing in Kubernetes](#)
 - [Overview](#)
 - [Building Kubernetes and Running the Tests](#)
 - [Cleaning up](#)
 - [Advanced testing](#)
 - [Extracting a specific version of Kubernetes](#)
 - [Bringing up a cluster for testing](#)
 - [Debugging clusters](#)
 - [Debugging an E2E test with a debugger \(delve \)](#)
 - [Local clusters](#)
 - [Testing against local clusters](#)
 - [Version-skewed and upgrade testing](#)
 - [Test jobs naming convention](#)
 - [Kinds of tests](#)
 - [Viper configuration and hierarchical test parameters.](#)
 - [Conformance tests](#)
 - [Continuous Integration](#)



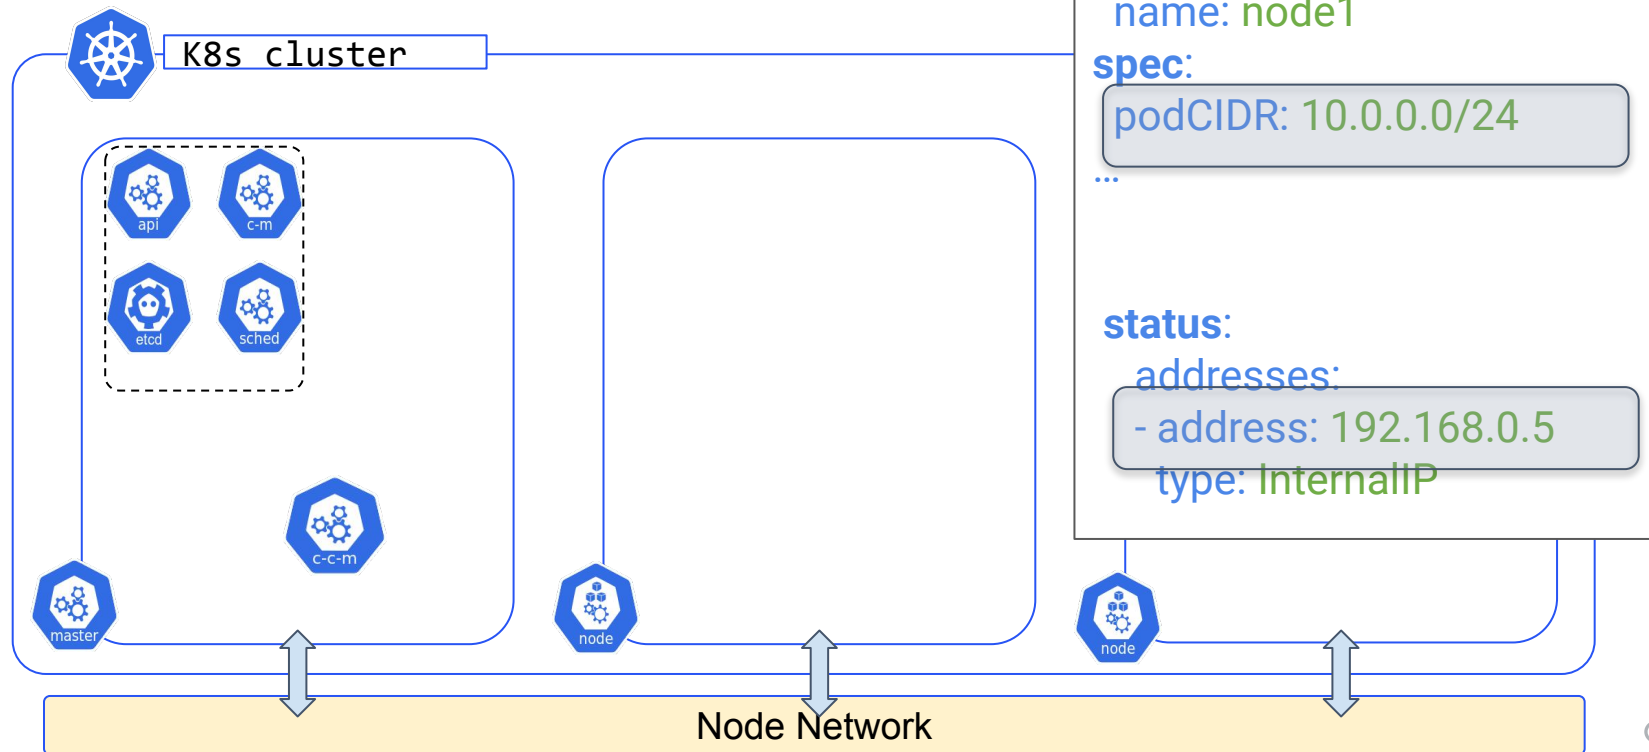
Kubernetes Implementation: one of multiple



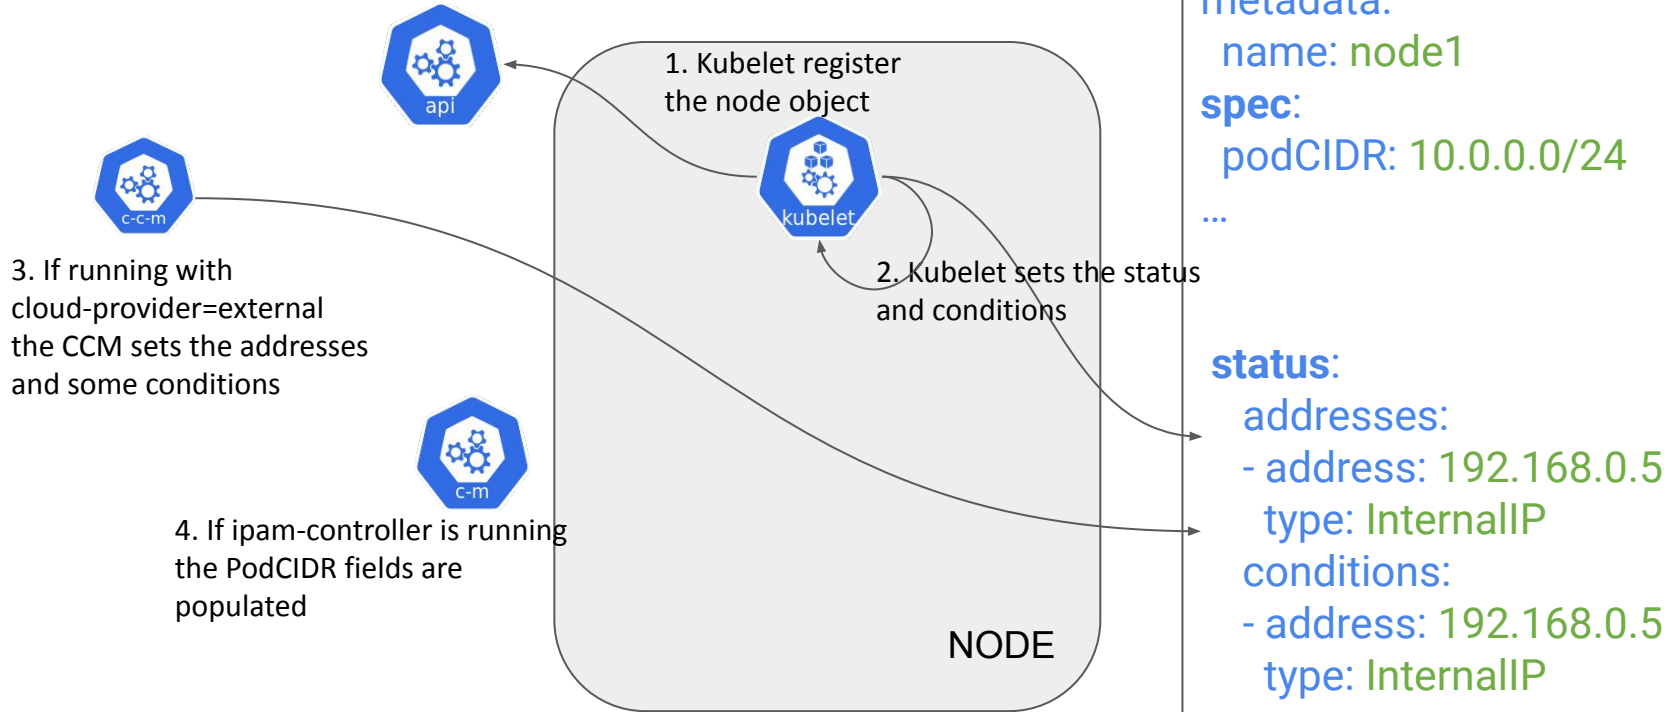
Kubernetes Networking: end to end principle



Kubernetes networking: Nodes

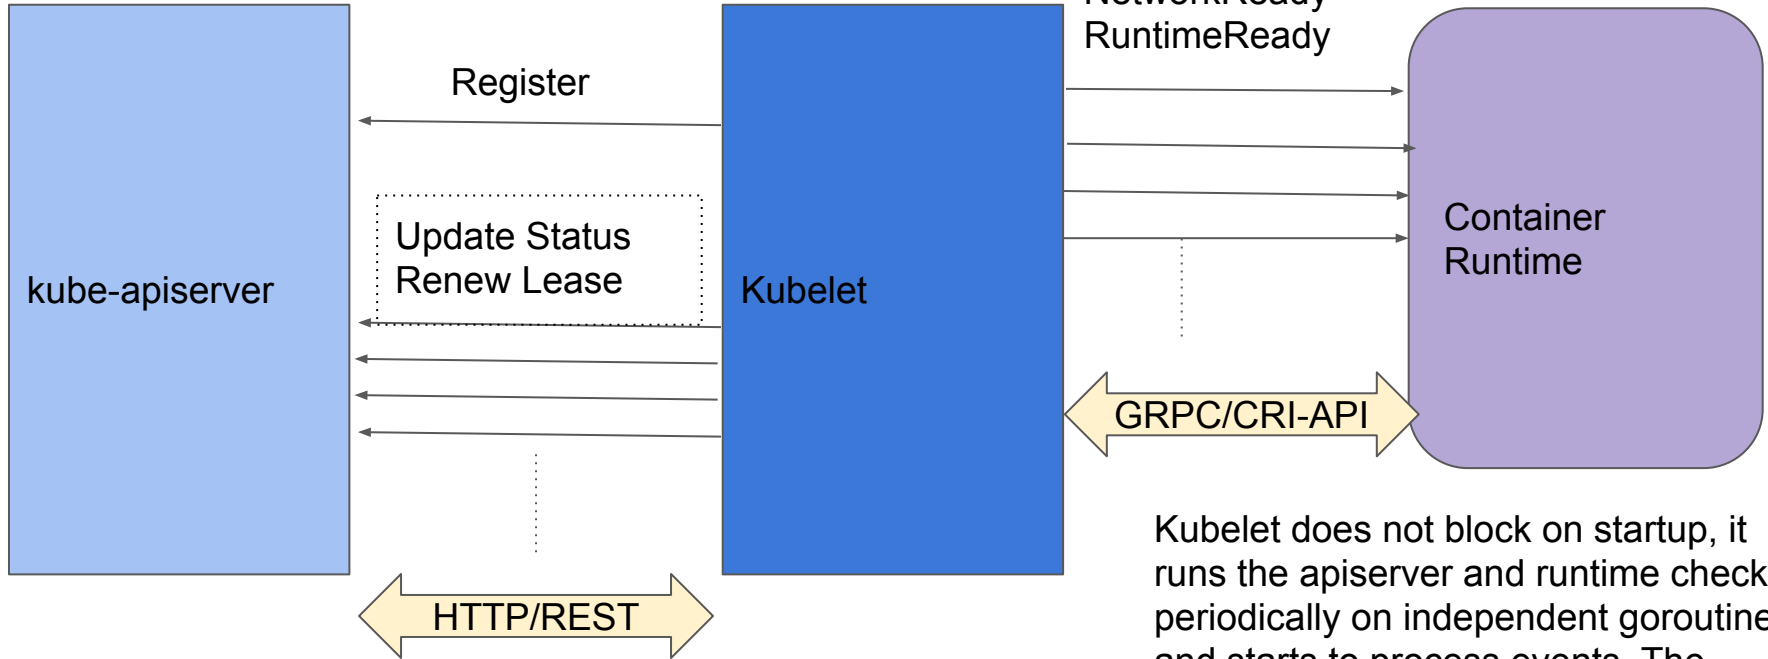


Kubernetes networking: Nodes



Node Initialization

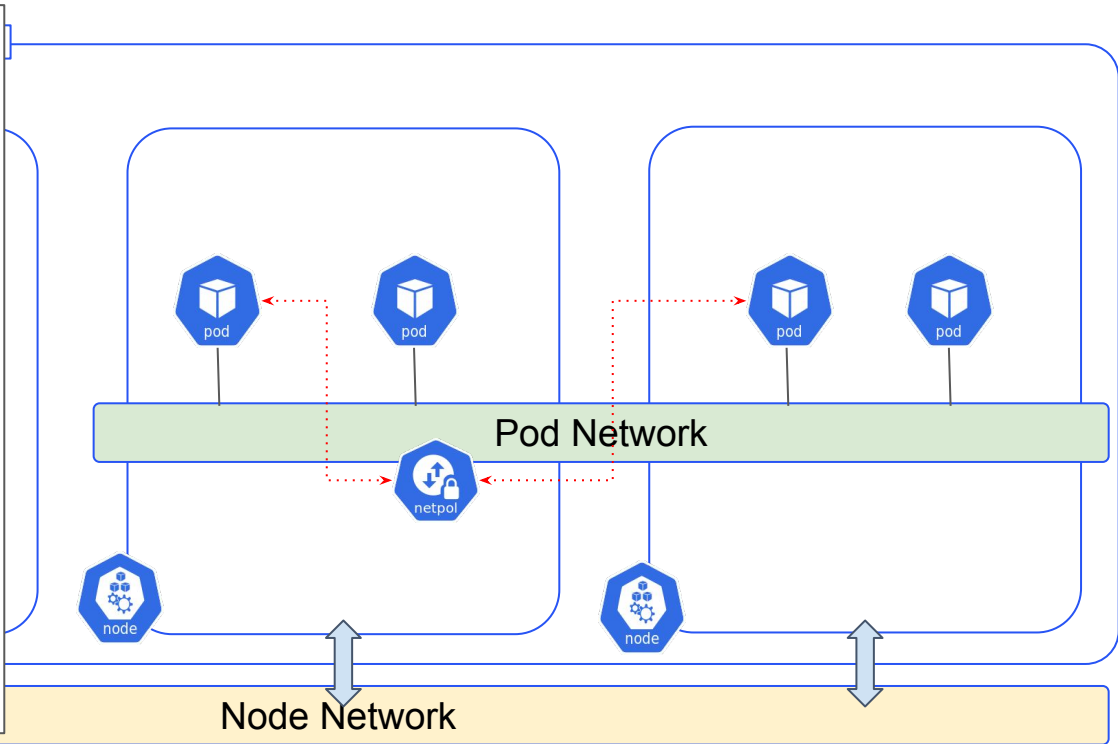
One of the runtime checks is “NetworkReady”, this is implemented on the container runtimes just as a check that a configuration file for the CNI exists



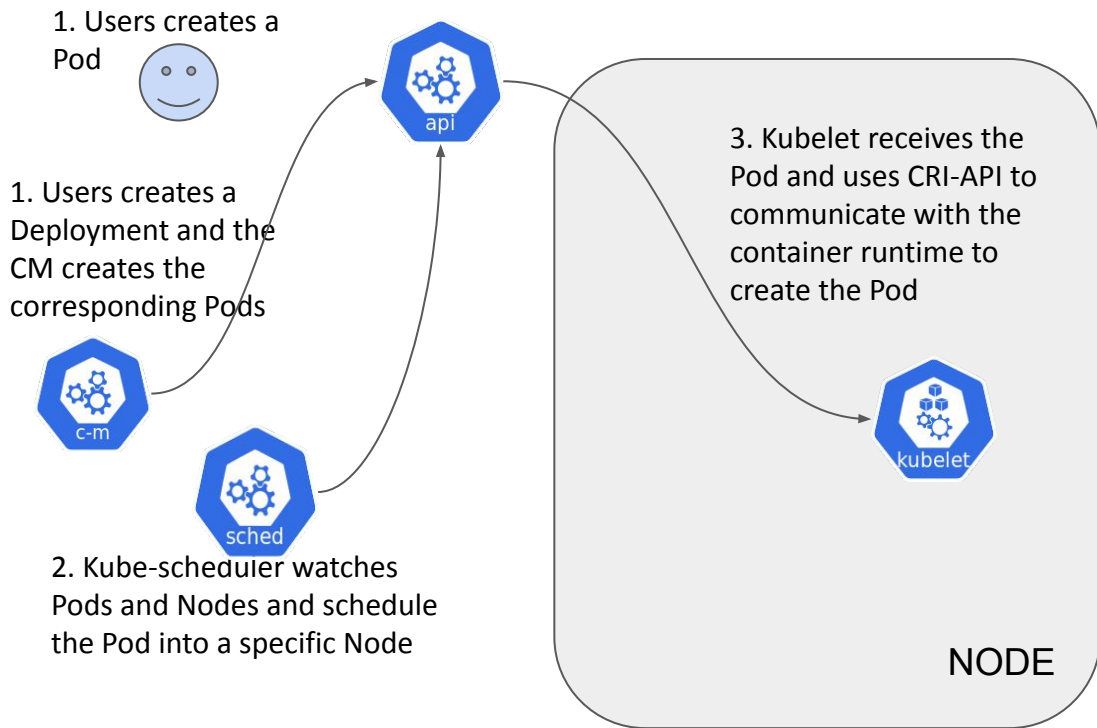
Kubelet does not block on startup, it runs the apiserver and runtime checks periodically on independent goroutines and starts to process events. The Node is not declared Ready until all the checks are OK

Kubernetes networking: Pods

```
apiVersion: v1
kind: Pod
metadata:
  name: app-http
  namespace: default
spec:
  containers:
  - name: app
    image: myapp:0.1
  ...
status:
  phase: Running
  podIP: 10.0.0.12
```



Kubernetes networking: Pods

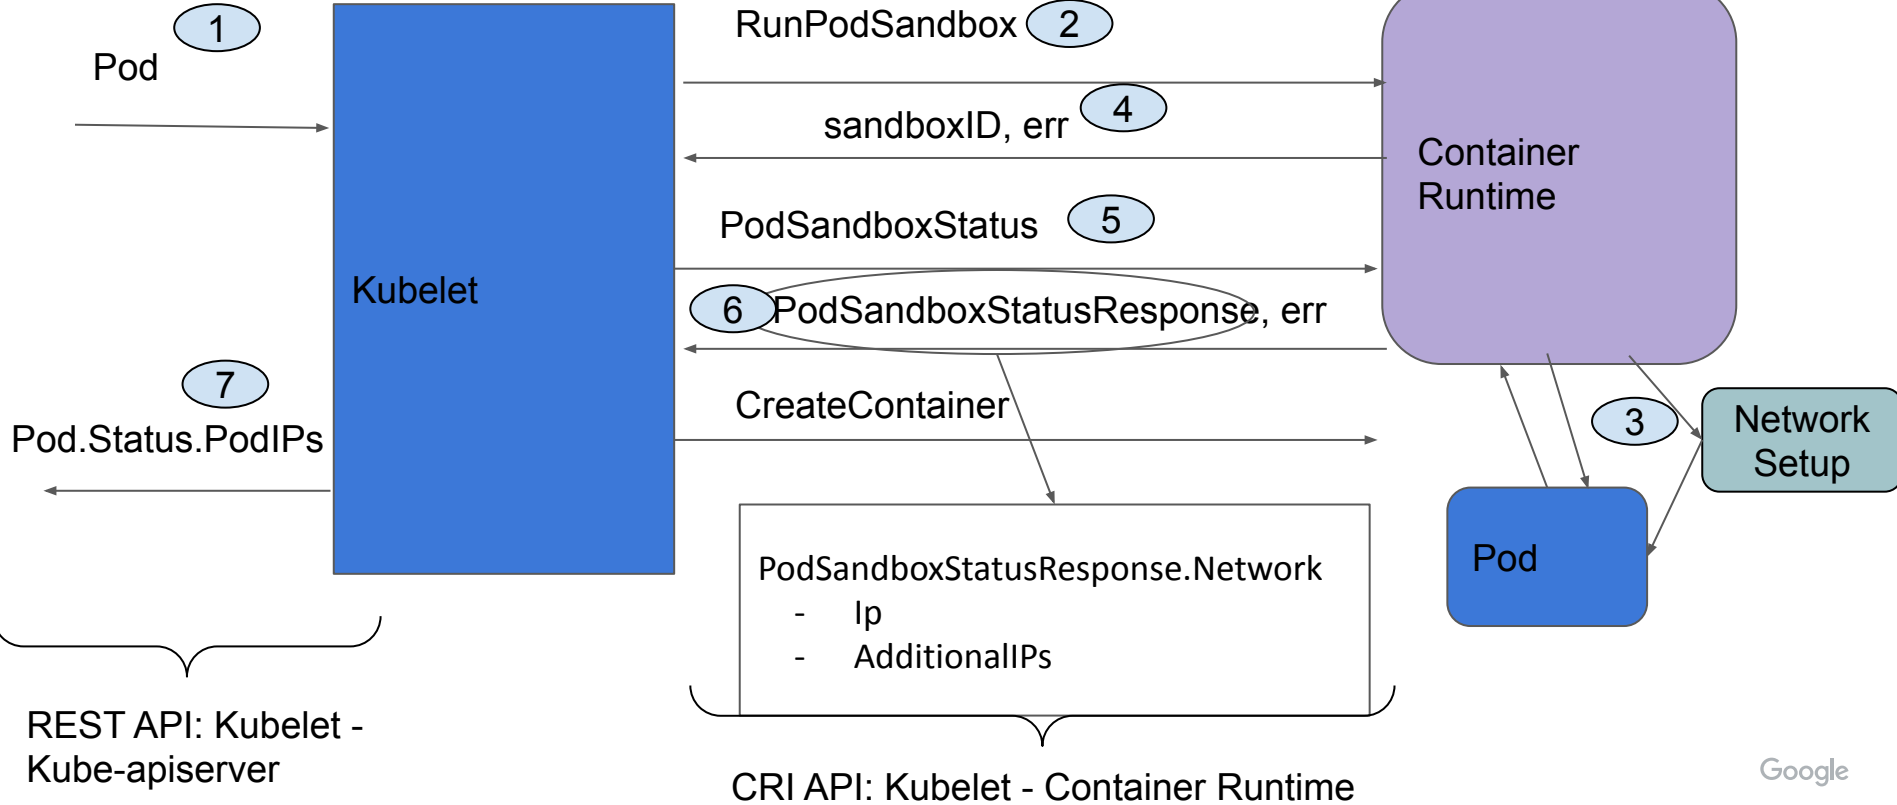


```
apiVersion: v1
kind: Pod
metadata:
  name: app-http
  namespace: default
spec:
  containers:
  - name: app
    image: myapp:0.1
  ...
status:
  phase: Running
  podIP: 10.0.0.12
```

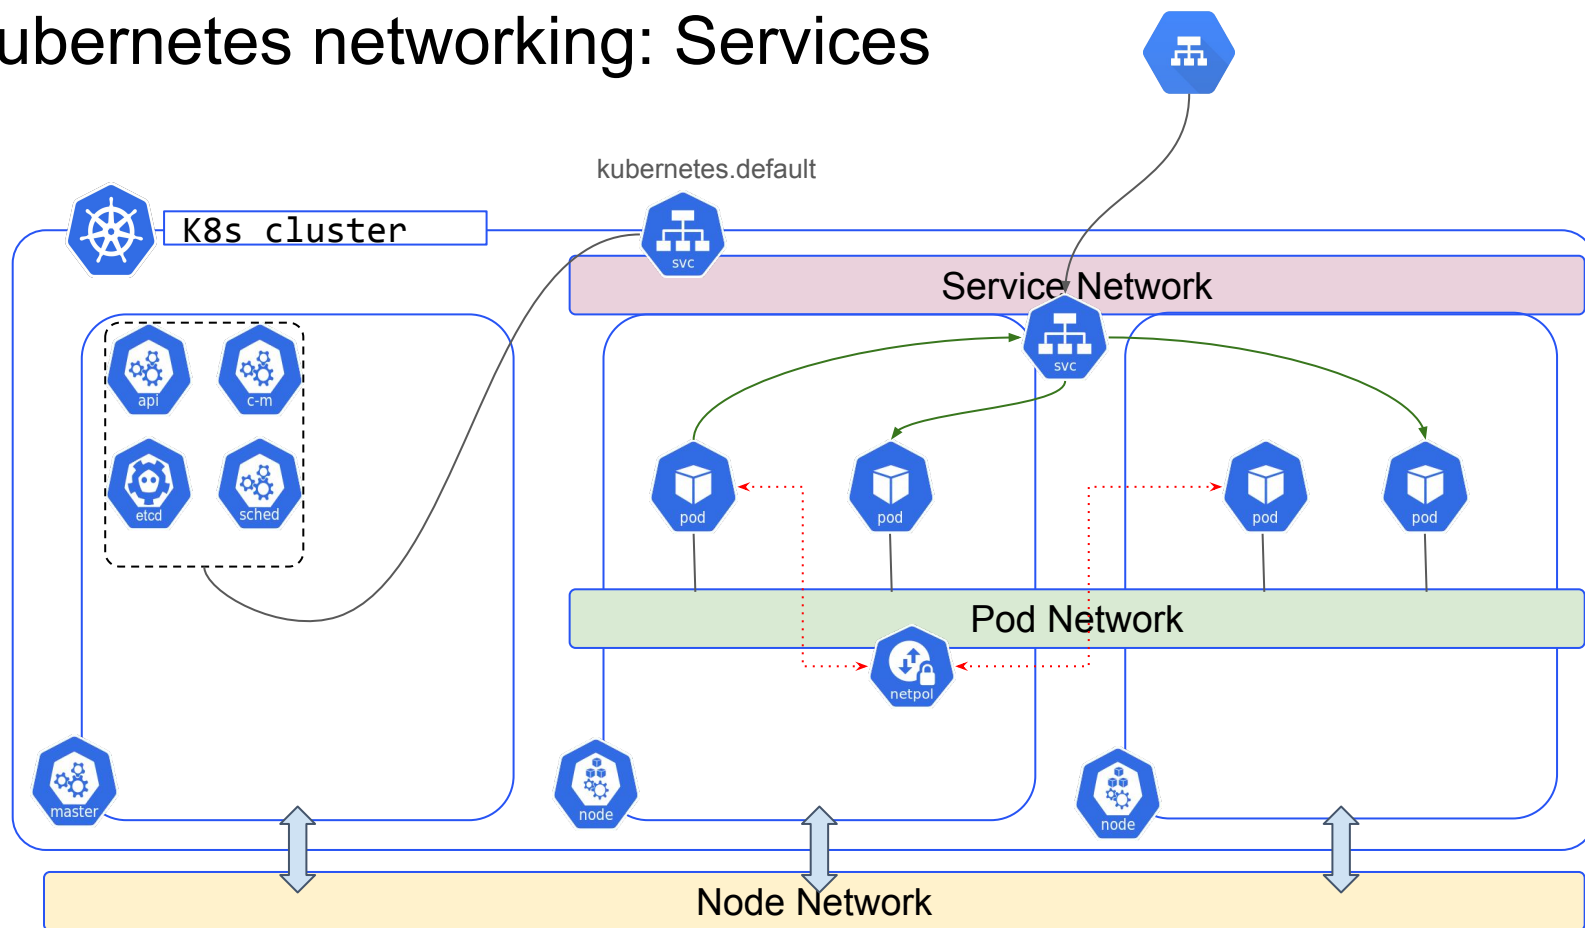
Kubernetes networking: CRI

RunPodSandboxRequest.PodSandboxConfig

- DNSConfig
- PortMapping
- Hostname



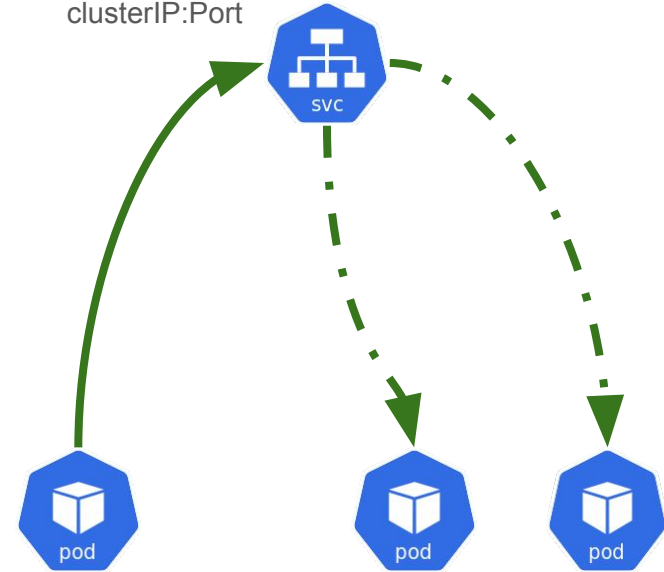
Kubernetes networking: Services



Kubernetes Services: ClusterIP

```
apiVersion: v1
kind: Service
metadata:
  name: service
spec:
  clusterIP: 10.96.0.1
  clusterIPs:
  - 10.96.0.1
  internalTrafficPolicy: ClusterIP
  ipFamilies:
  - IPv4
  ipFamilyPolicy: SingleStack
  ports:
  - name: https
    port: 443
    protocol: TCP
    targetPort: 6443
  type: ClusterIP
```

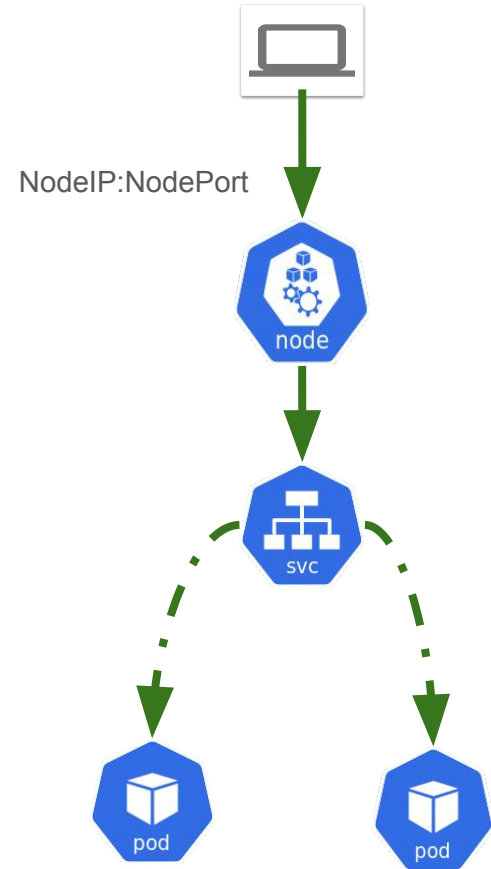
```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: server-deployment
  labels:
    app: MyApp
spec:
  replicas: 2
  selector:
    matchLabels:
      app: MyApp
  template:
    metadata:
      labels:
        app: MyApp
    spec:
      terminationGracePeriodSeconds: 30
      containers:
      - name: agnhost
        image:
          k8s.gcr.io/e2e-test-images/agnhost:2.39
        args:
          - netexec
          - --http-port=80
```



Kubernetes Services: NodePort

```
apiVersion: v1
kind: Service
metadata:
  name: service
spec:
  clusterIP: 10.96.0.1
  clusterIPs:
  - 10.96.0.1
  internalTrafficPolicy: Cluster
  ipFamilies:
  - IPv4
  ipFamilyPolicy: SingleStack
  ports:
  - name: https
    port: 443
    protocol: TCP
    targetPort: 6443
    nodePort: 31023
    type: NodePort
```


```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: server-deployment
labels:
  app: MyApp
spec:
  replicas: 2
  selector:
    matchLabels:
      app: MyApp
  template:
    metadata:
      labels:
        app: MyApp
    spec:
      terminationGracePeriodSeconds: 30
      containers:
      - name: agnhost
        image:
          k8s.gcr.io/e2e-test-images/agnhost:2.39
        args:
          - netexec
          - --http-port=80
```



Kubernetes Services: Load Balancer

```
apiVersion: v1
kind: Service
metadata:
  name: service
spec:
  clusterIP: 10.96.0.1
  clusterIPs:
  - 10.96.0.1
  internalTrafficPolicy: Cluster
  ipFamilies:
  - IPv4
  ipFamilyPolicy: SingleStack
  ports:
  - name: https
    port: 443
    protocol: TCP
    targetPort: 6443
  type: LoadBalancer
status:
  loadBalancer:
    ingress:
    - ip: 202.34.23.12
```

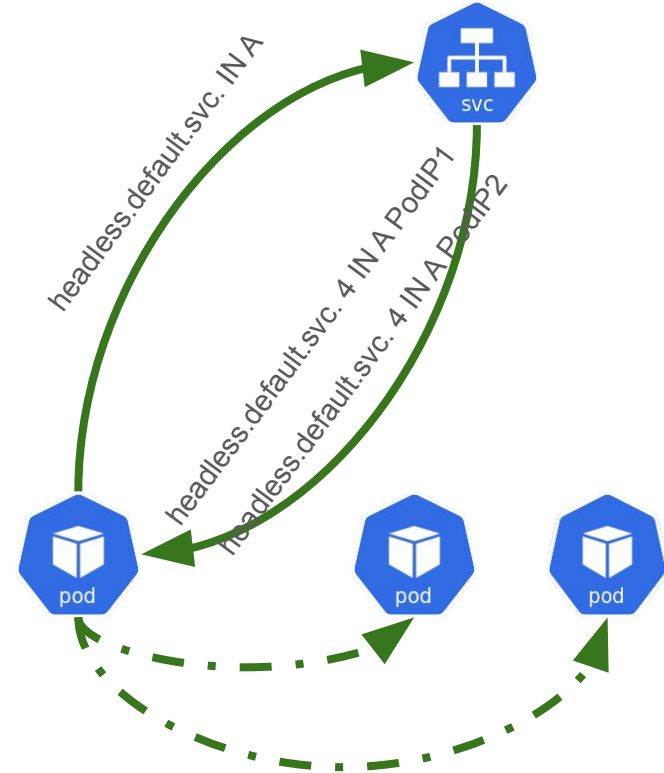
```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: server-deployment
labels:
  app: MyApp
spec:
  replicas: 2
  selector:
    matchLabels:
      app: MyApp
  template:
    metadata:
      labels:
        app: MyApp
    spec:
      terminationGracePeriodSeconds: 30
      containers:
      - name: agnhost
        image:
          k8s.gcr.io/e2e-test-images/agnhost:2.39
        args:
        - netexec
        - --http-port=80
```



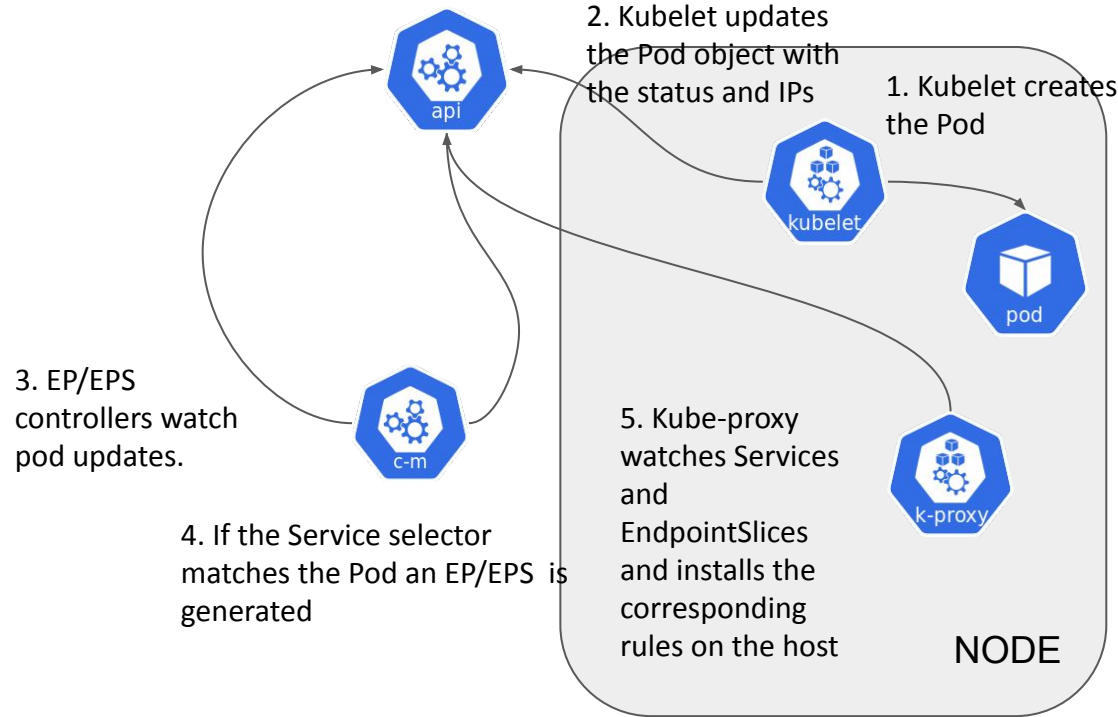
Kubernetes Services: Headless

```
apiVersion: v1
kind: Service
metadata:
  name: headless
spec:
  clusterIP: None
  clusterIPs:
  - None
  internalTrafficPolicy: Cluster
  ipFamilies:
  - IPv4
  ipFamilyPolicy: SingleStack
  ports:
  - name: https
    port: 443
    protocol: TCP
    targetPort: 6443
  type: ClusterIP
```

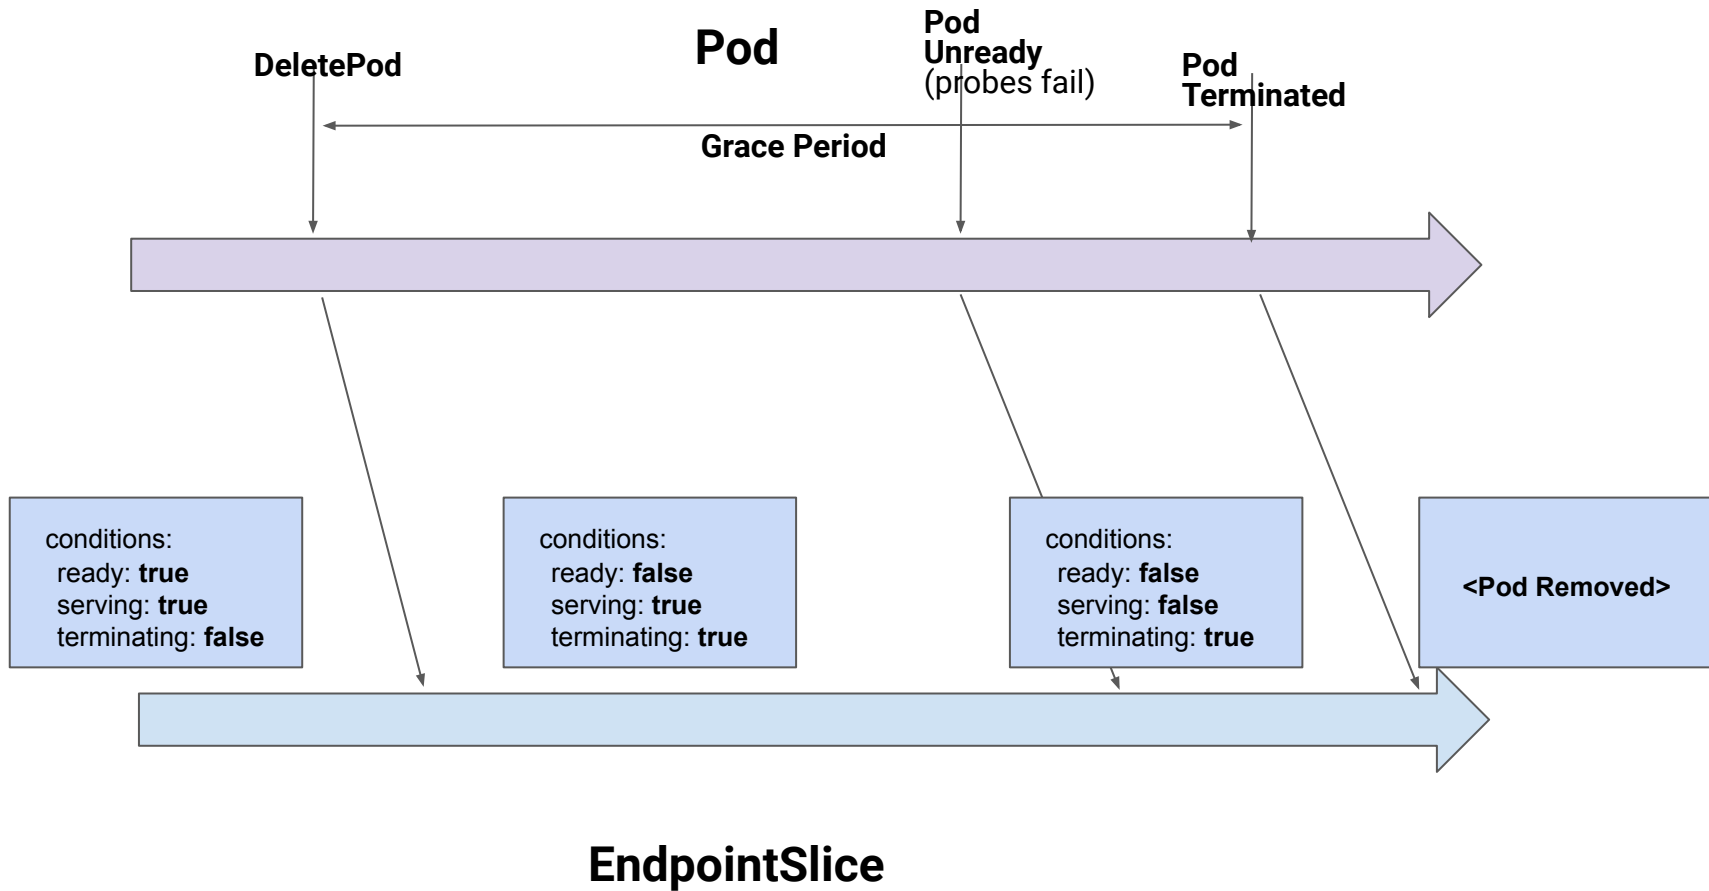
```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: server-deployment
labels:
  app: MyApp
spec:
  replicas: 2
  selector:
    matchLabels:
      app: MyApp
  template:
    metadata:
      labels:
        app: MyApp
    spec:
      terminationGracePeriodSeconds: 30
      containers:
      - name: agnhost
        image:
          k8s.gcr.io/e2e-test-images/agnhost:2.39
        args:
          - netexec
          - --http-port=80
```

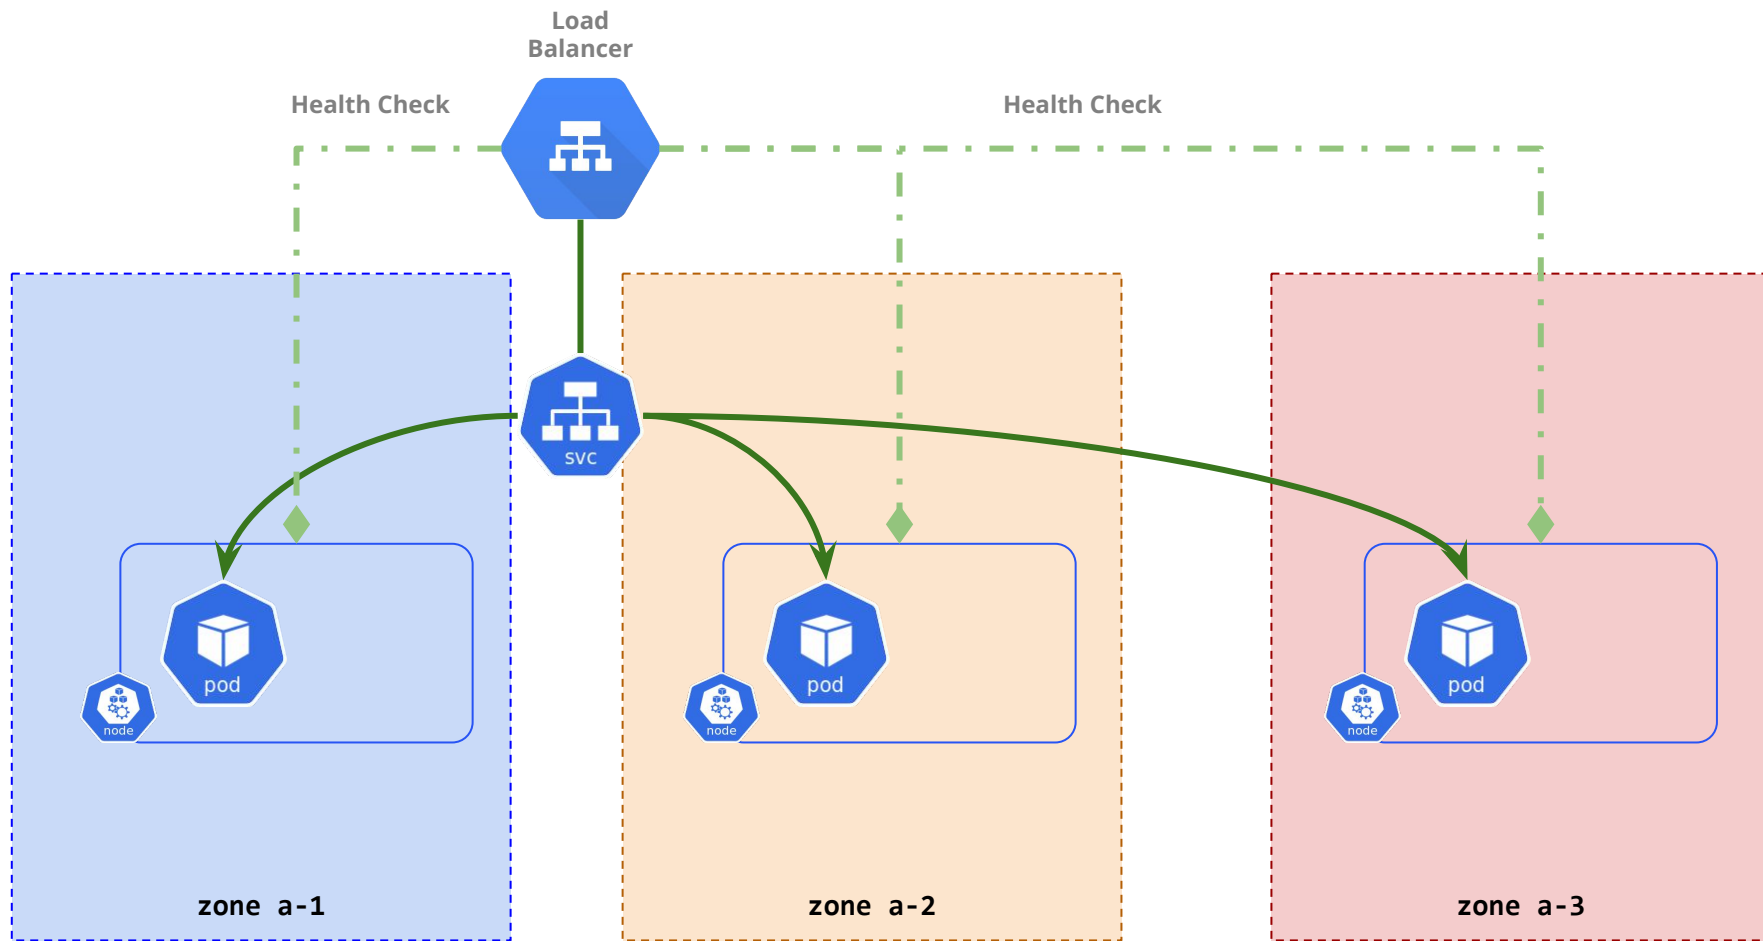


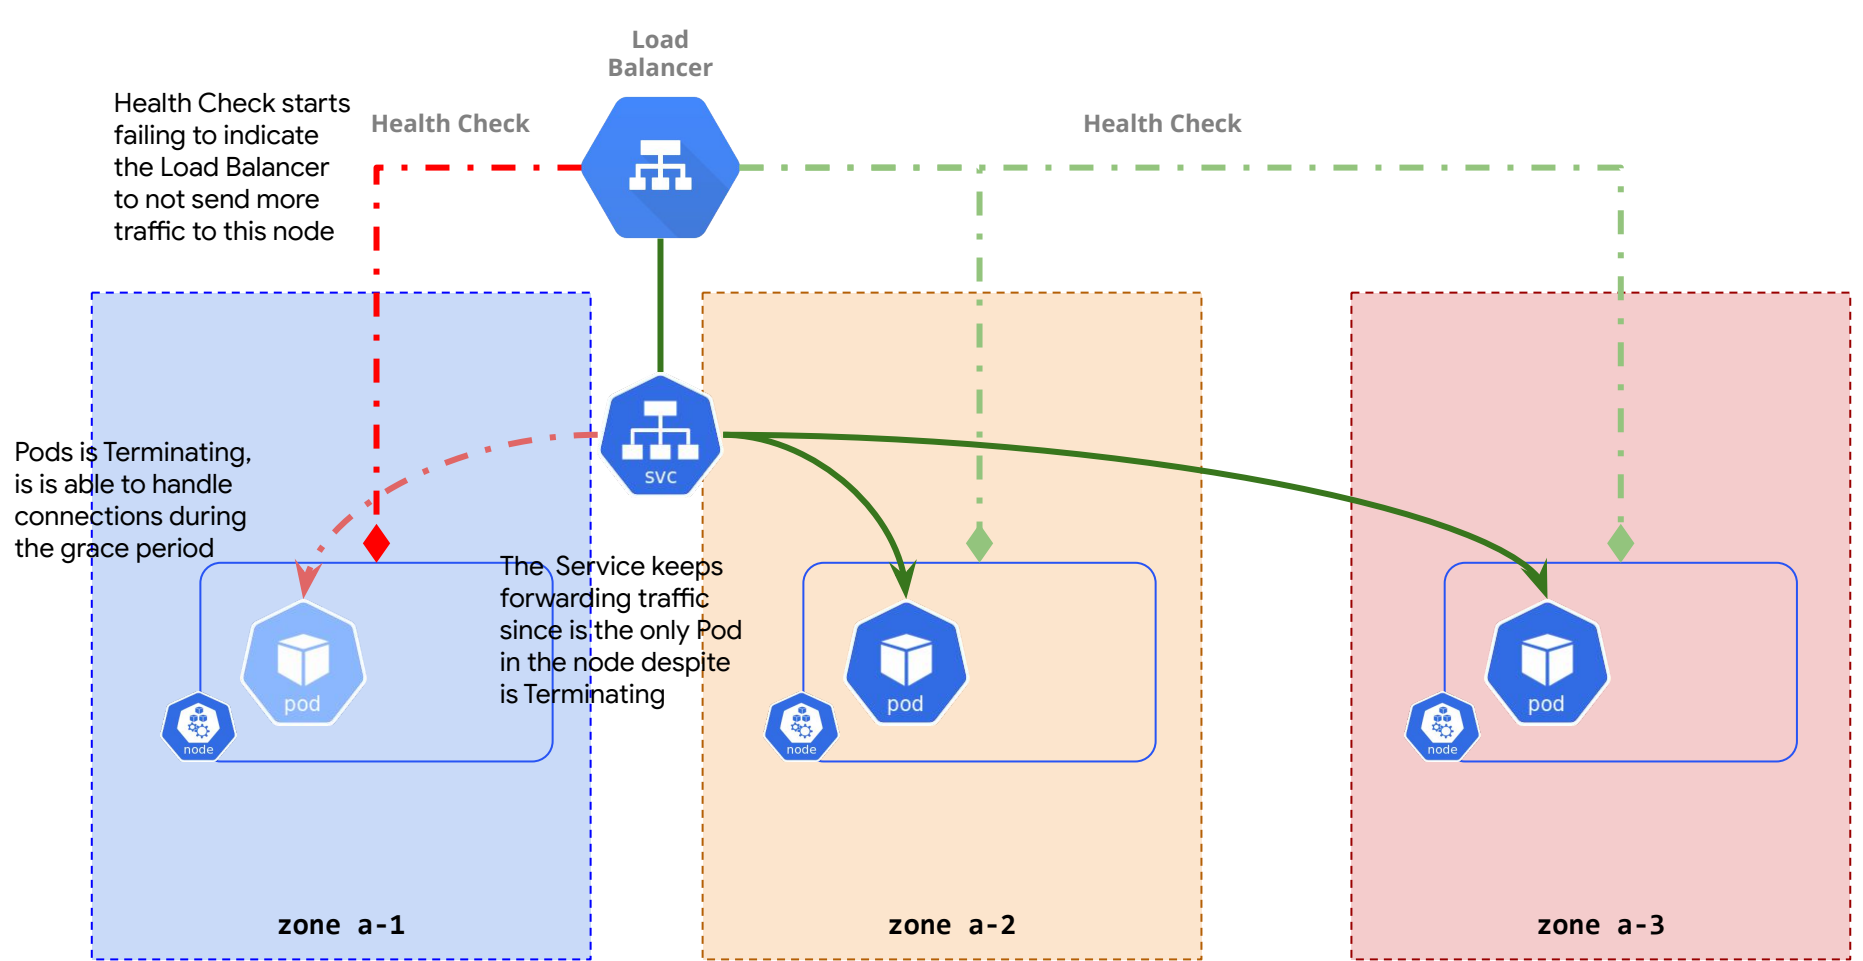
Kubernetes Services

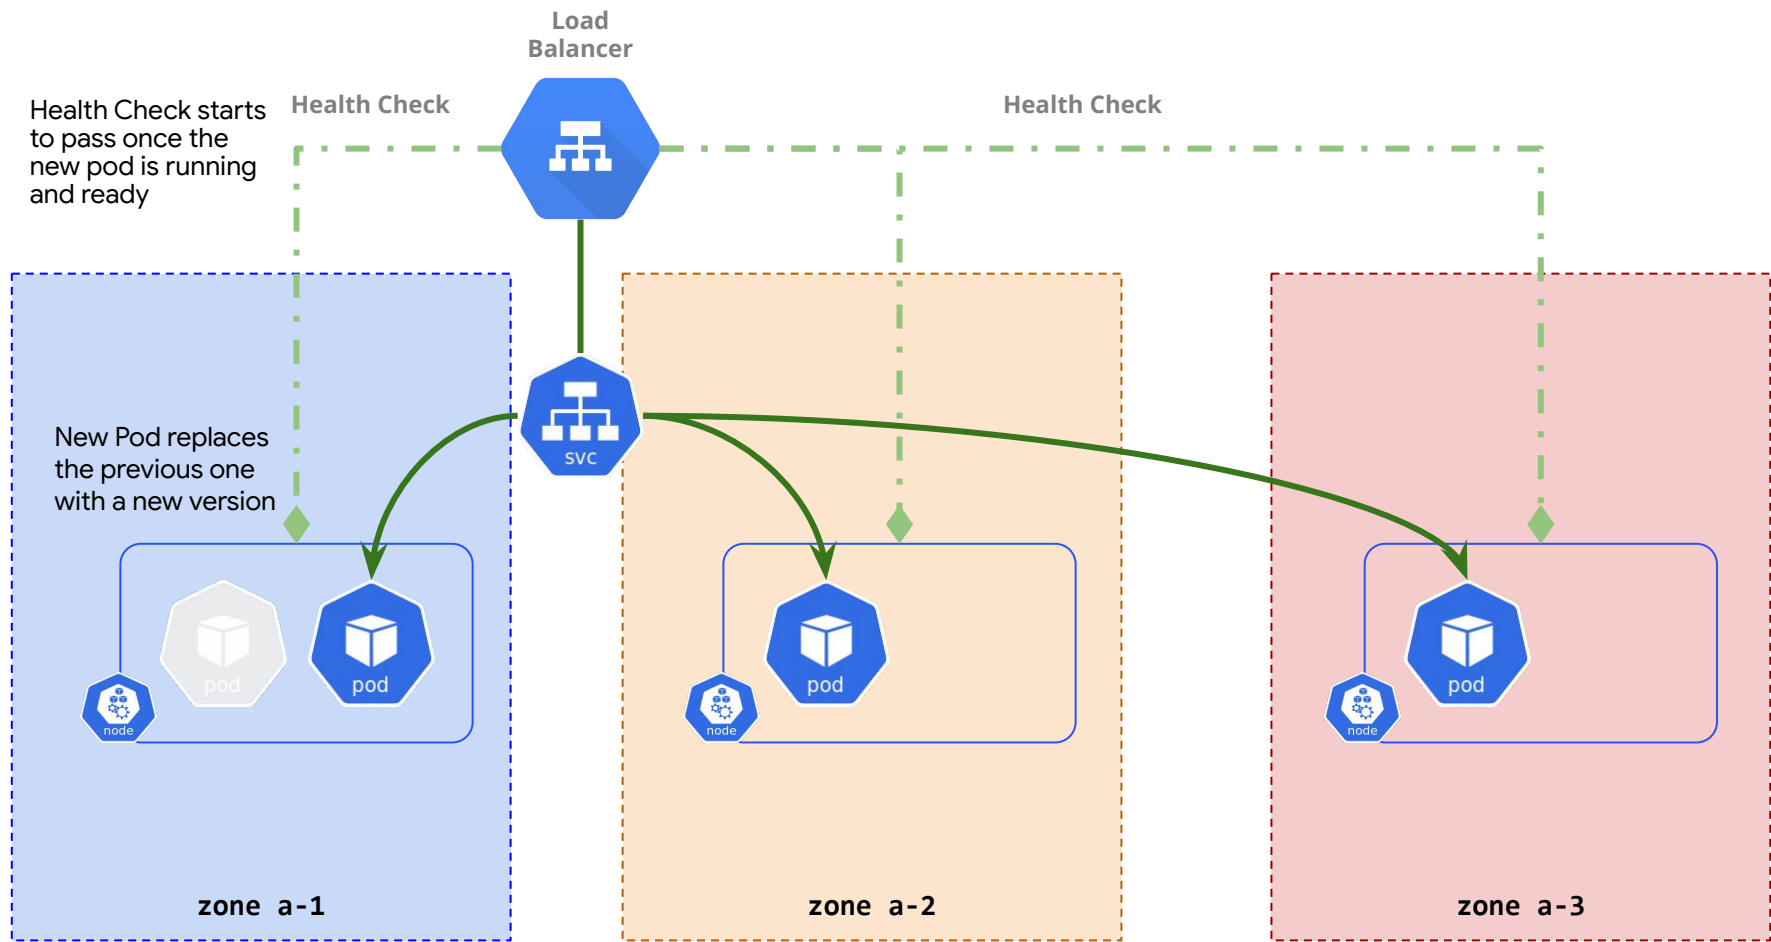


```
apiVersion: v1
kind: Service
metadata:
  name: lb-app
  namespace: default
spec:
  clusterIP: 10.96.0.13
  selector:
    app: MyApp
  type: ClusterIP
  ports:
    - name: port8080
      port: 8080
      protocol: TCP
      targetPort: 8080
```



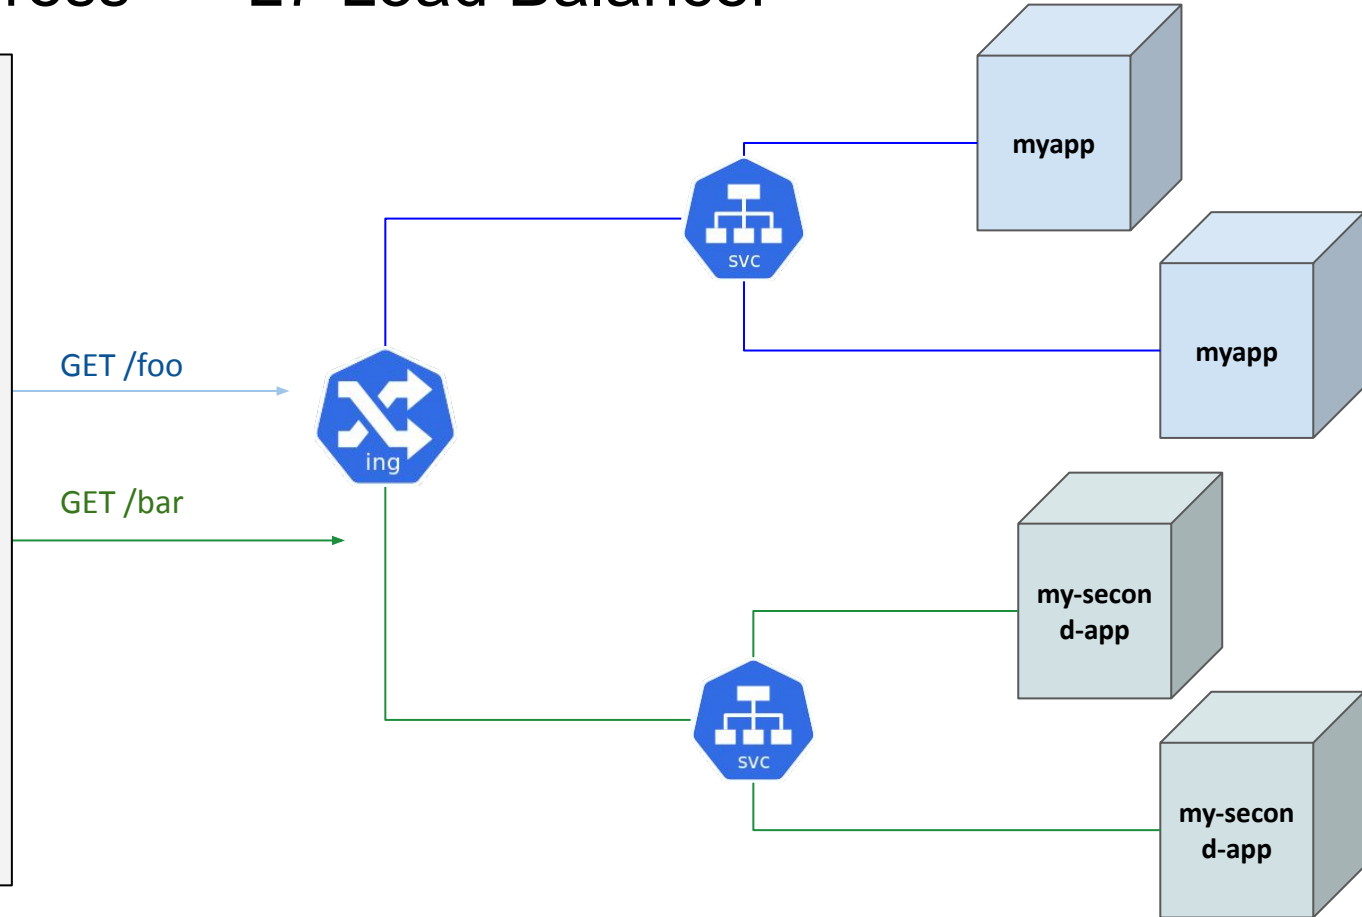




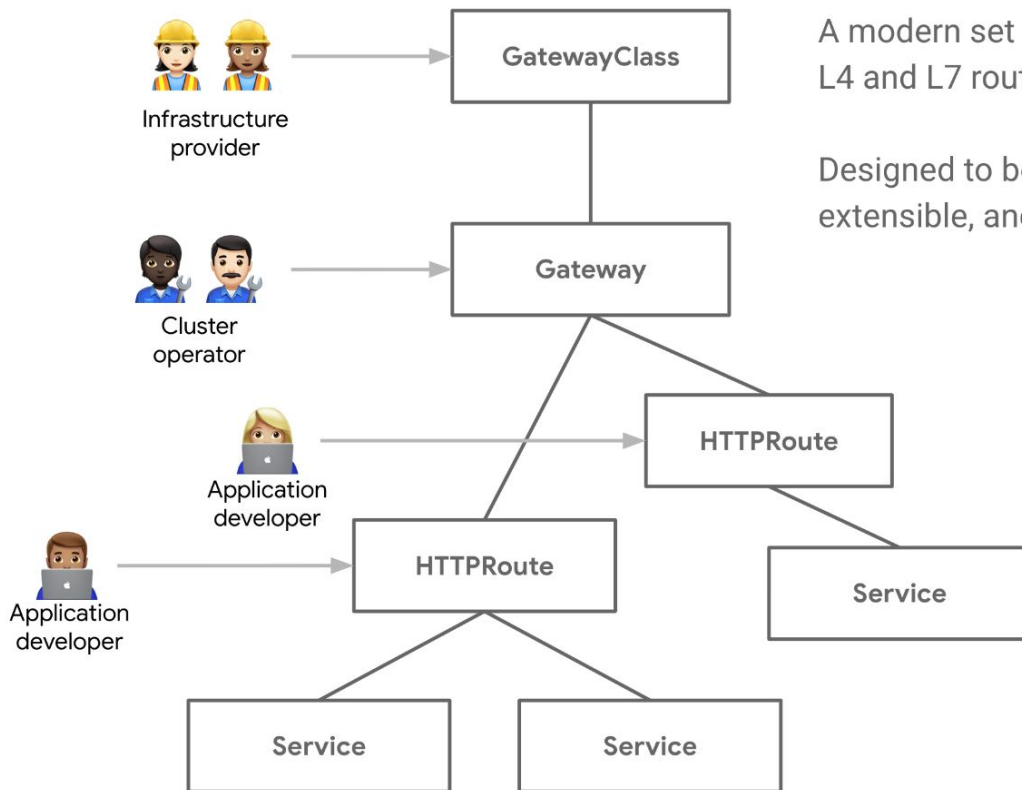


Kubernetes Ingress == L7 Load Balancer

```
apiVersion: v1
kind: Ingress
metadata:
  name: minimal-ingress
spec:
  ingressClassName: nginx-example
  rules:
  - http:
    paths:
    - path: /foo
      pathType: Prefix
      backend:
        service:
          name: myapp
          port:
            number: 80
    - path: /bar
      pathType: Prefix
      backend:
        service:
          name: my-second-app
          port:
            number: 80
```

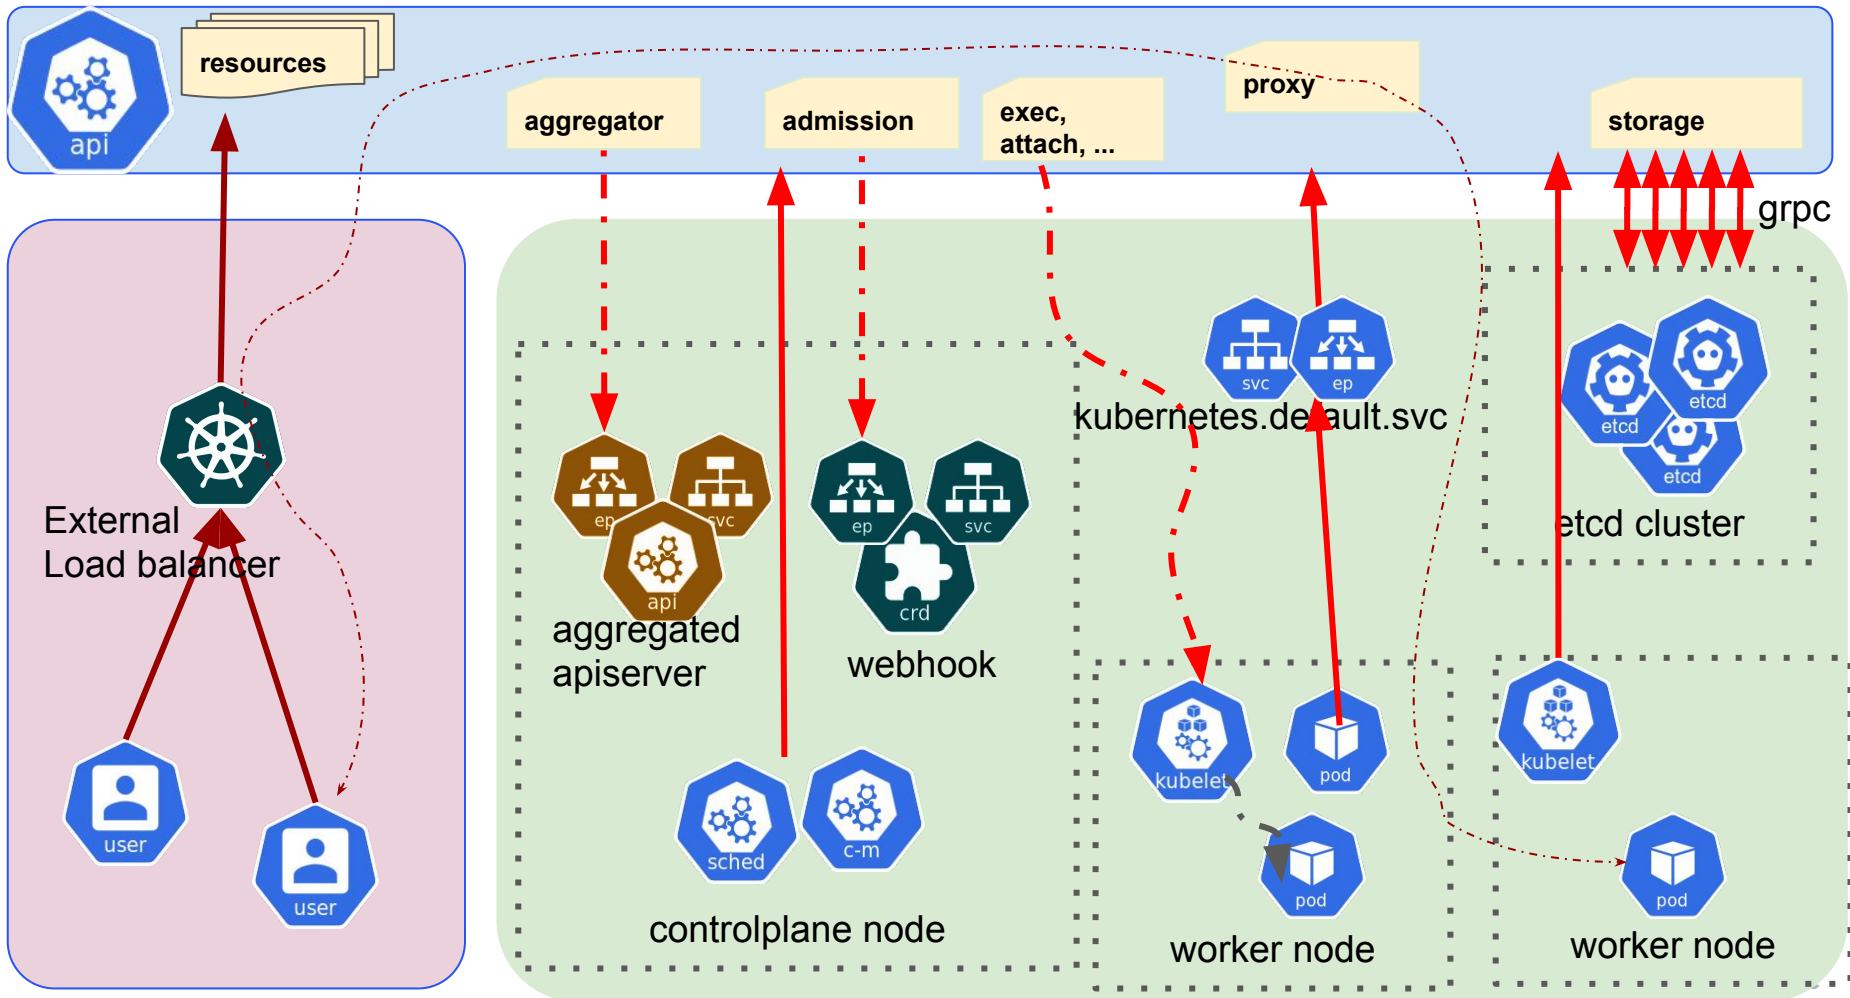


Gateway API aims to standardize the space



A modern set of APIs for deploying L4 and L7 routing in Kubernetes

Designed to be generic, expressive, extensible, and role-oriented



What is next

The screenshot shows a GitHub Kanban board for "WIP: SIG Network KEPs". The board is organized into five columns representing different stages of development: New, Pre-Alpha, Alpha, Beta, and GA. Each column contains a list of items, each with a title and an enhancement ID. The board also includes a search bar, a filter by keyword or by field, and a "Save" button.

WIP: SIG Network KEPs

Filter by keyword or by field

New 1

- enhancements #4427
Relaxed DNS search string validation

Pre-Alpha 3

- enhancements #3698
Multi-Network
- enhancements #784
Graduate the kube-proxy ComponentConfig to v1beta1
- enhancements #4410
[KNI] Kubernetes Network reimagined (interface)

Alpha 5

- enhancements #1860
Make Kubernetes aware of the LoadBalancer behaviour
- enhancements #1880
Multiple Service CIDRs
- enhancements #3836
Kube-proxy improved ingress connectivity reliability
- enhancements #3866
ntables kube-proxy backend
- enhancements #4004
Deprecate status.nodeInfo.kubeProxyVersion field

Beta 4

- enhancements #2681
Field status.hostIPs added for Pod
- enhancements #3458
Remove transient node predicates from KCCM's service controller
- enhancements #3705
Cloud Dual-Stack --node-ip Handling
- enhancements #2433
Topology Aware Routing

GA 1

- enhancements #3668
Reserve nodeport ranges for dynamic and static allocation