

AusNOG 7th September 2023 : Kubernetes autoscaling and load balancing

Yin Loong CHAO
Akamai



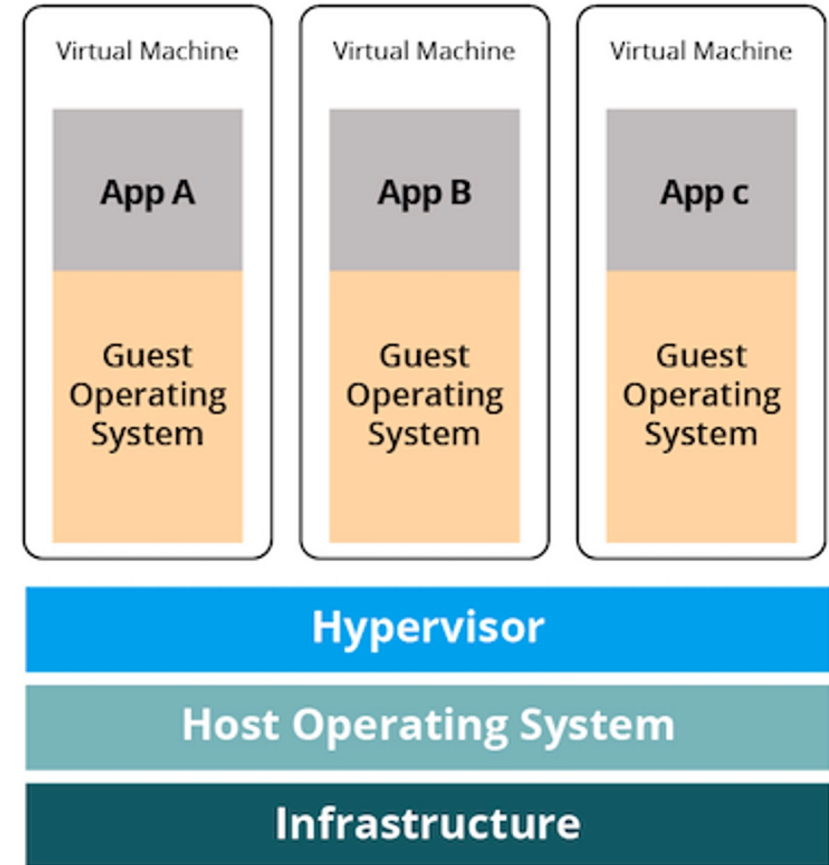
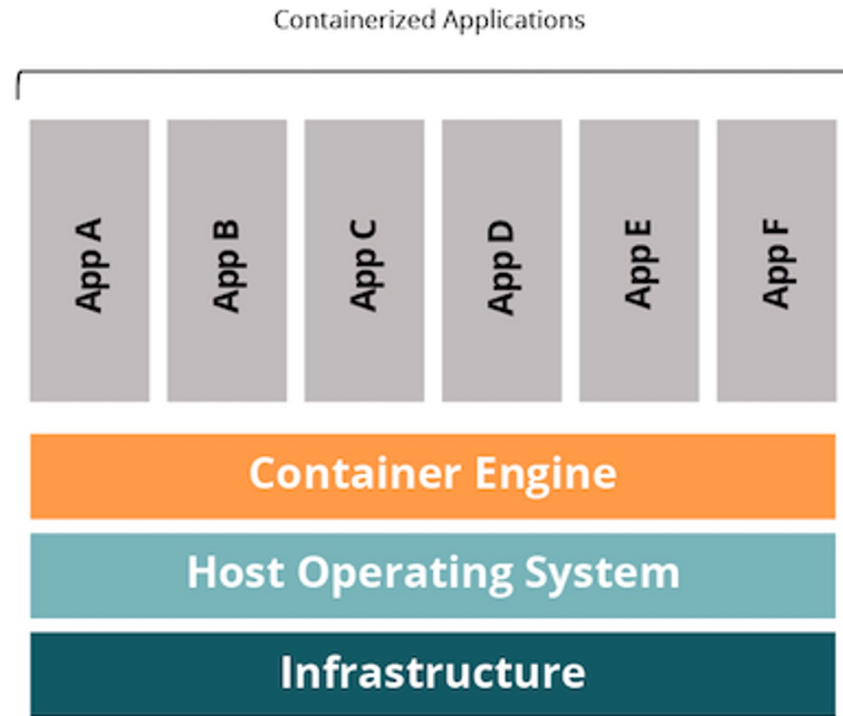
Purpose of talk

- Kubernetes and containers defacto cloud operating system
- Network operators increasingly using K8S for their workloads like netops, AIOps, 5G functions
- Take advantage of the native scaling flexibility and load balancing capabilities of kubernetes

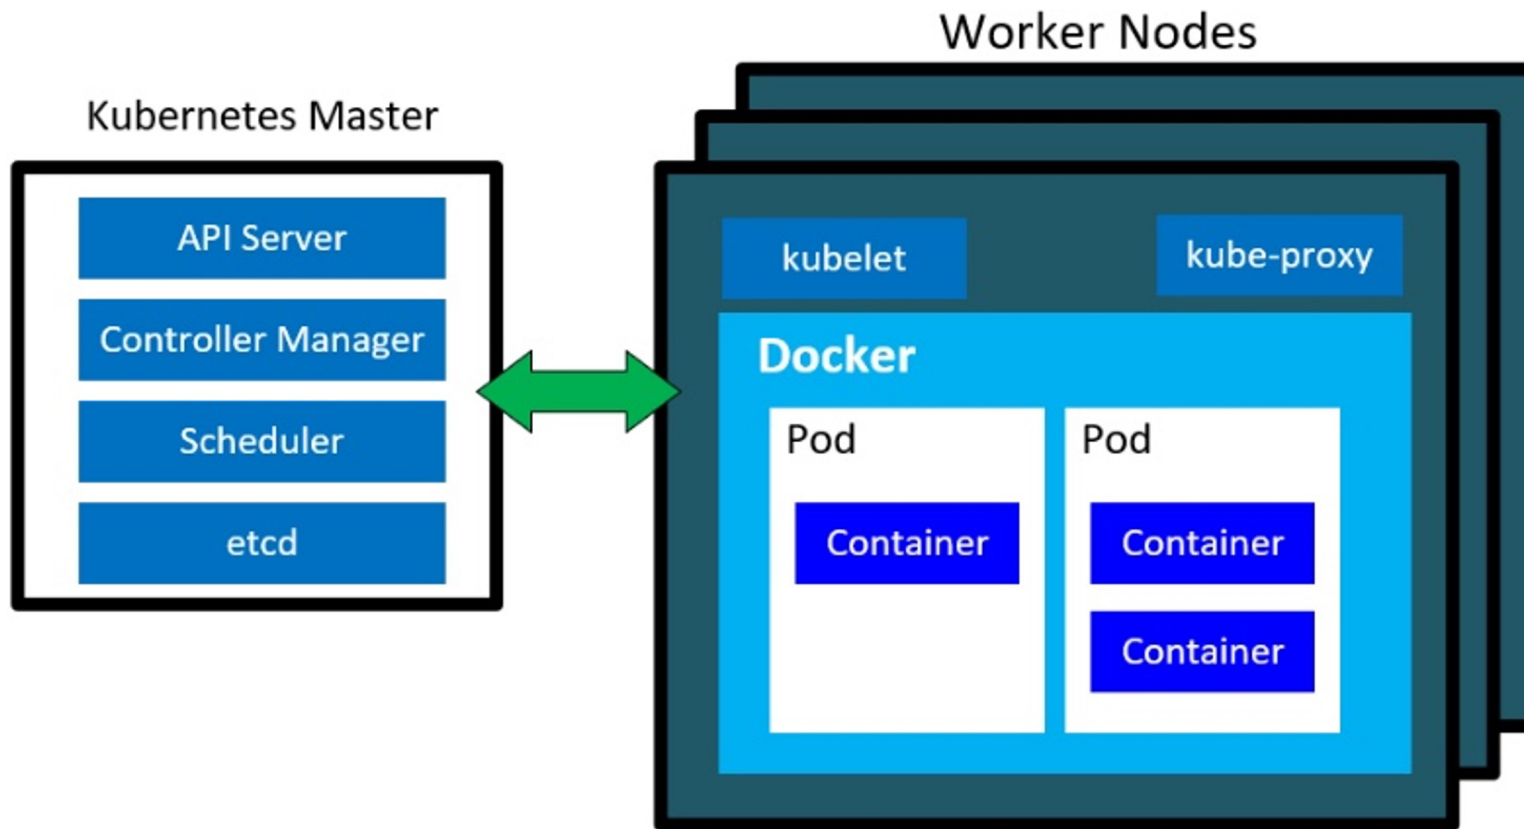
Agenda

- Kubernetes Intro
- Setting up pod autoscaling + node auto scaling.
- Load balancing in Kubernetes pods
- Use cases

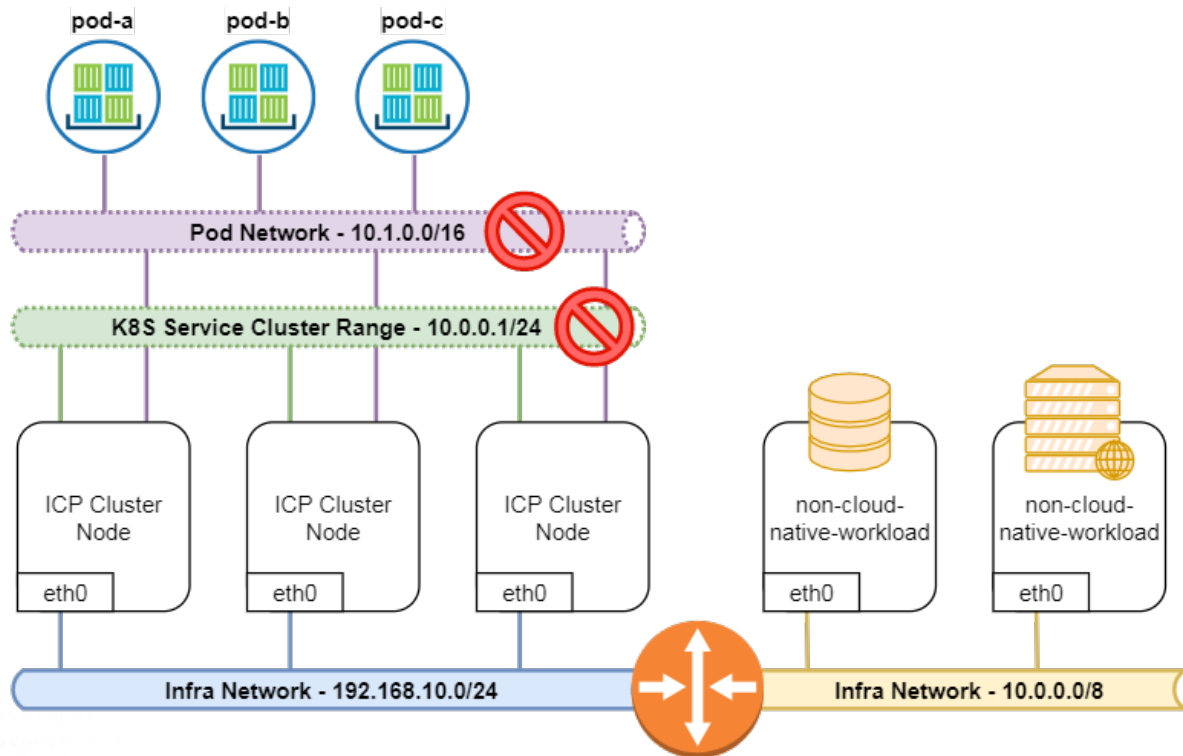
Containers Intro



Kubernetes



Kubernetes networking model

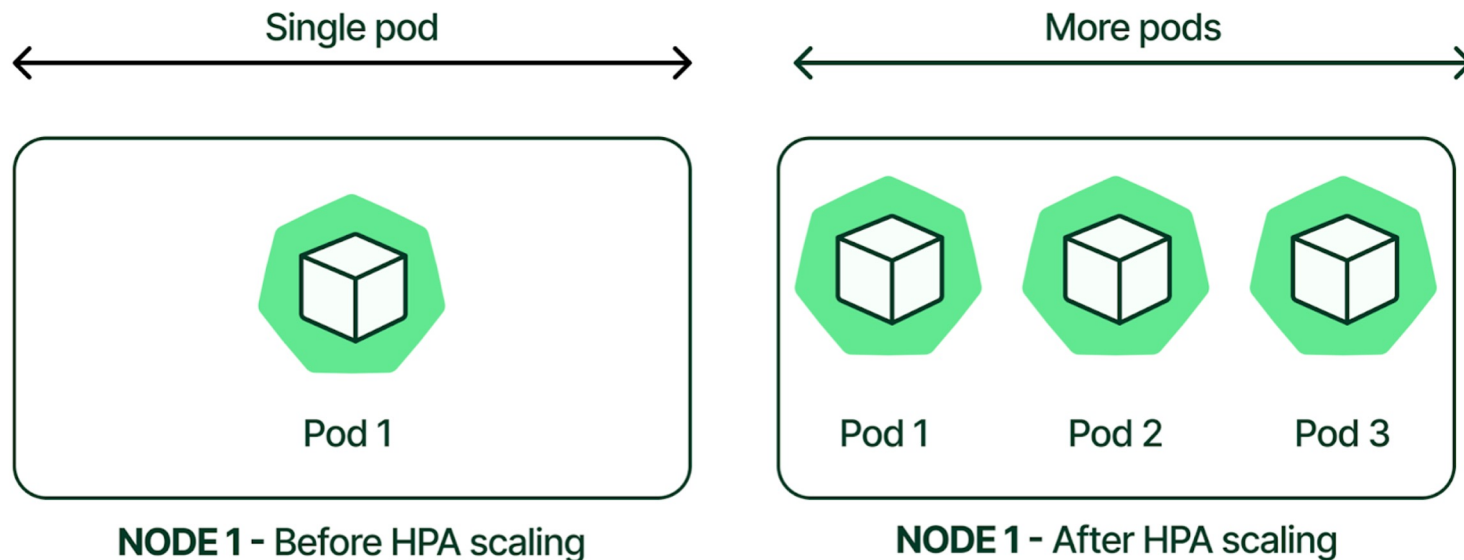


Kubernetes capacity planning

- Cluster Level
- Resource Level
- Node Level

Kubernetes Horizontal Pod Autoscaler Intro

HPA Autoscaling



Installing metrics server

- Install metrics server
- Once you install metrics server you can use top command

```
[(base) ychao@sin-mpdlh k8shpa % kubectl top pods
```

NAME	CPU(cores)	MEMORY(bytes)
linode-kubernetes-6c5f6b644d-599w8	0m	21Mi
php-apache-5b56f9df94-kw7f6	1m	11Mi

```
(base) ychao@sin-mpdlh k8shpa %
```

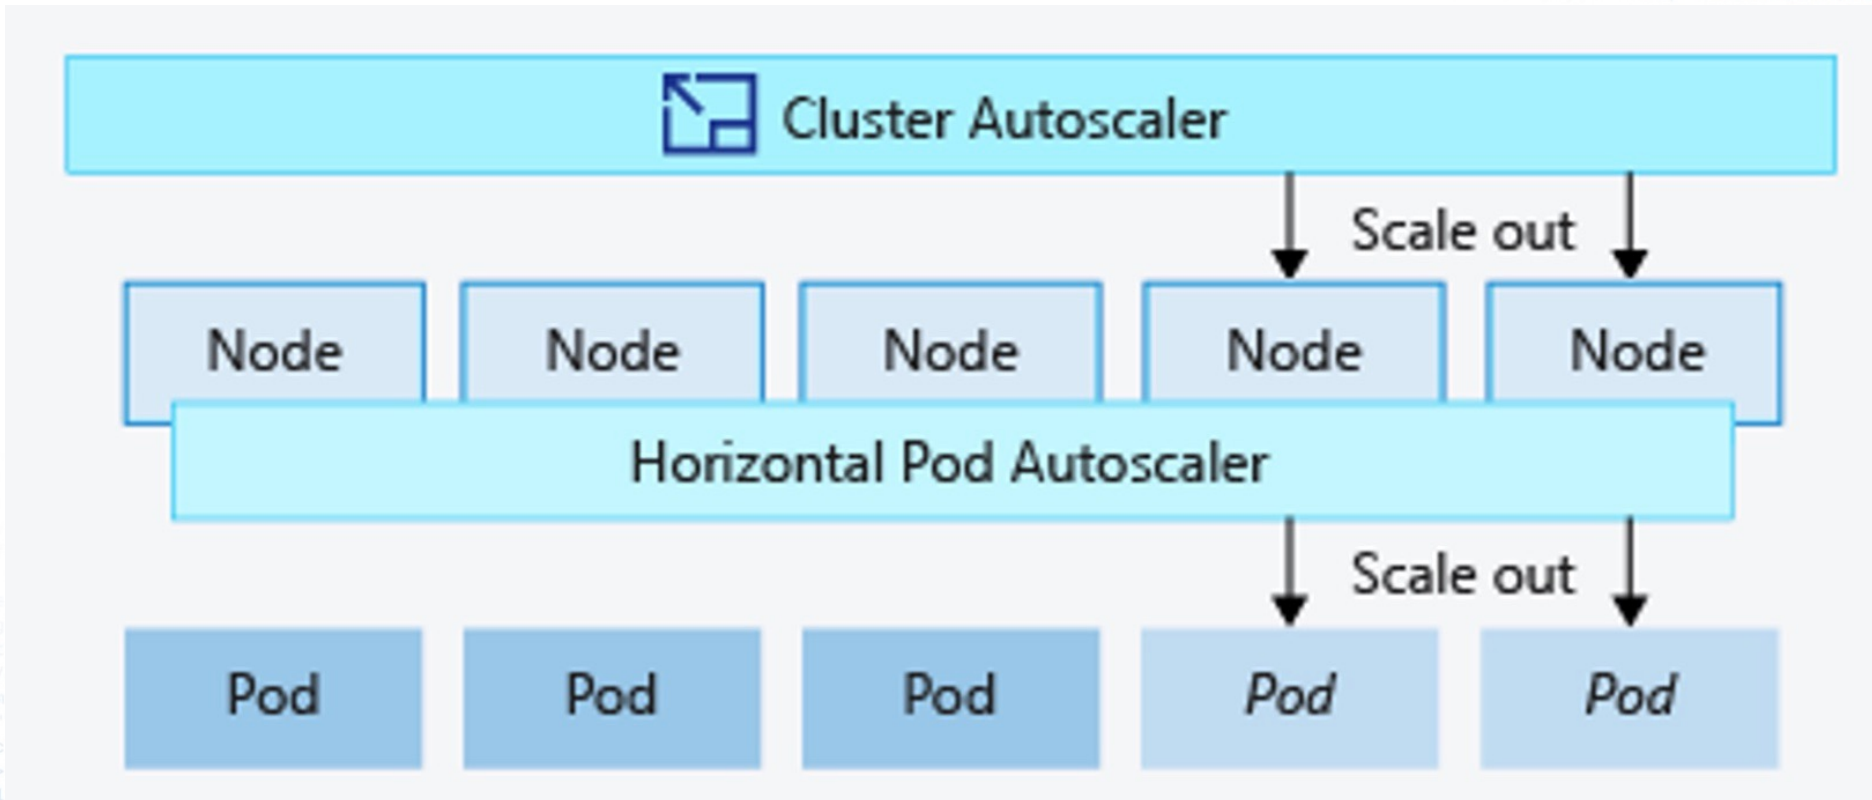
Install Horizontal Pod Autoscaler Operator

- Install horizontal pod autoscaler operator for the application deployment

```
((base) ychao@sin-mpdlh k8shpa % kubectl autoscale deployment linode-kubernetes --cpu-percent=10 --min=1 --max=10
Error from server (AlreadyExists): horizontalpodautoscalers autoscaling "linode-kubernetes" already exists
```

```
Error from server (AlreadyExists): horizontalpodautoscalers
(base) ychao@sin-mpdlh k8shpa % kubectl get hpa -o yaml
apiVersion: v1
items:
- apiVersion: autoscaling/v2
  kind: HorizontalPodAutoscaler
  metadata:
    creationTimestamp: "2023-01-11T10:28:40Z"
  managedFields:
  - apiVersion: autoscaling/v1
    fieldsType: FieldsV1
    fieldsV1:
      f:spec:
        f:maxReplicas: {}
        f:minReplicas: {}
        f:scaleTargetRef: {}
        f:targetCPUUtilizationPercentage: {}
    manager: kubectl-autoscale
    operation: Update
    time: "2023-01-11T10:28:40Z"
    apiVersion: autoscaling/v2
```

Kubernetes Node Autoscaler



Enable node autoscaler

The screenshot displays the 'Node Pools' management interface. A modal dialog titled 'Autoscale Pool' is open, allowing configuration for a specific node pool. The dialog includes a checkbox for 'Autoscaler' (checked), and input fields for 'Min' (3) and 'Max' (3) node constraints. The background shows a list of node pools under the heading 'Dedicated 4GB', with columns for 'Linode', 'Pool ID', and 'Pool Name'. Buttons for 'Recycle All Nodes', 'Add A Node Pool', 'Resize Pool', 'Recycle Pool Nodes', and 'Delete Pool' are visible in the interface.


Kubernetes Pods are ephemeral

- Each Pod has its own IP address
- IP address changes as Pods are destroyed and recreated

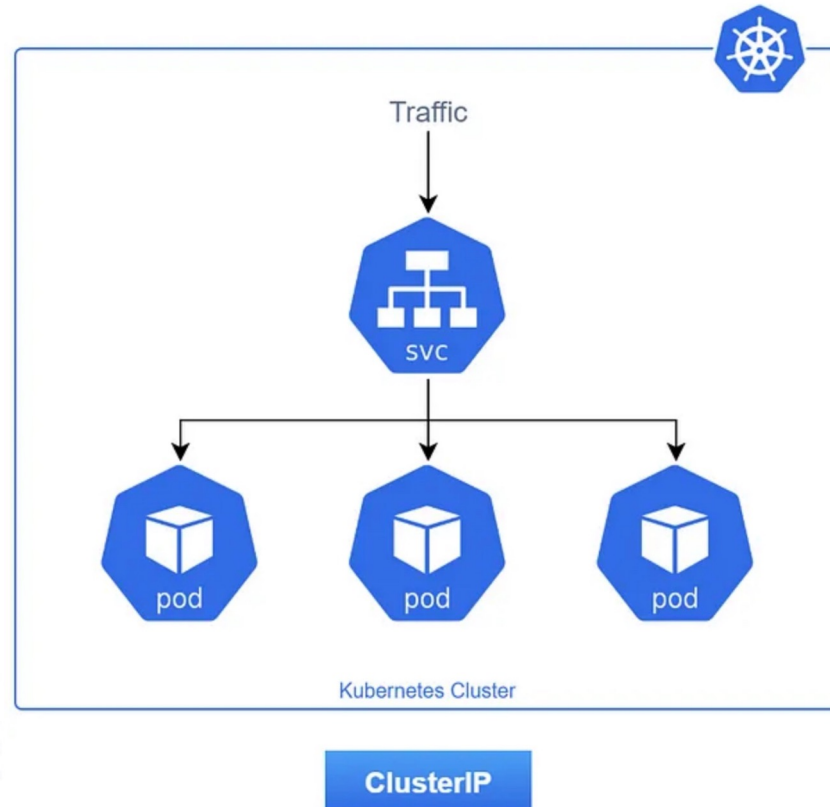


Kubernetes Service

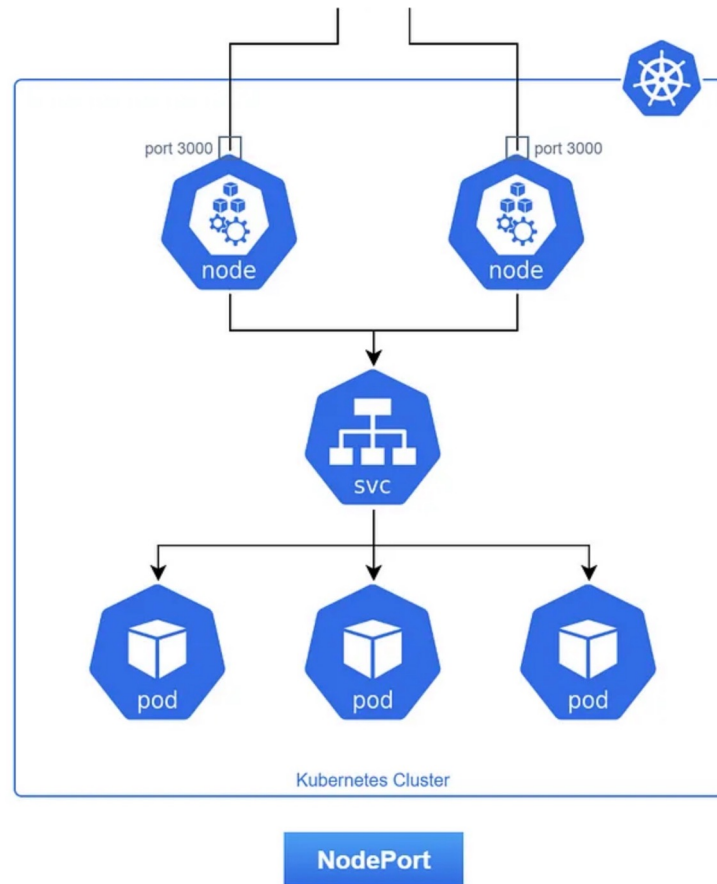
- Stable IP Address
- DNS name
- Load Balancing



Kubernetes Service Type - Cluster IP

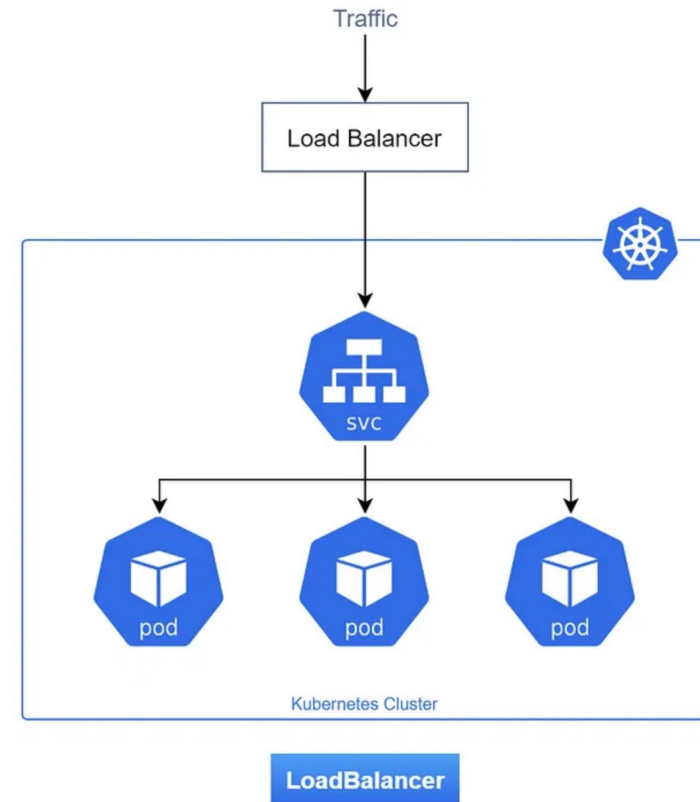


Kubernetes Service Type - Node Port



Kubernetes Service Type - Load Balancer

- For managed K8S like LKE, CSP will spin up a node balancer with an external IP
- Linode Cloud Controller Manager
<https://github.com/linode/linode-cloud-controller-manager>




Load Balancing within pods

- Kube-proxy updates iptables within nodes
- Iptables uses statistic method with probabilities

```
-A KUBE-SVC-GZJ05PIKZP44WMKS -m comment --comment "default/linode-kubernetes-service:http -> 10.2.0.8:3000" -m statistic --mode random --probability 0.33333333349 -j KUBE-SEP-TZ4FFJGK5UPMCD46
-A KUBE-SVC-GZJ05PIKZP44WMKS -m comment --comment "default/linode-kubernetes-service:http -> 10.2.1.5:3000" -m statistic --mode random --probability 0.50000000000 -j KUBE-SEP-T66P7EWYWFKISMU
-A KUBE-SVC-GZJ05PIKZP44WMKS -m comment --comment "default/linode-kubernetes-service:http -> 10.2.2.12:3000" -j KUBE-SEP-ESQKZSUP2CDSFK3J
```

Use case

- By using service Kubernetes resources, one can harness the power of Kubernetes for both scalability and self-healing enabling network operators to serve millions of customers at a very high level of availability while minimizing cost
- E.g. of workload for Network Operator, include any sort of IT workloads like websites, billing and even up to 5g Core Functions on Kubernetes.



(Anca Pavel, Fatih E. Nar, Federico Rossi, and Mathias Bogebrant, CC BY-SA 4.0)

Q&A



Containers Intro

