

# Hard core AKS

Alessandro Melchiori

@amelchiori

Founder & software developer @ CodicePlastico



# Let's start!

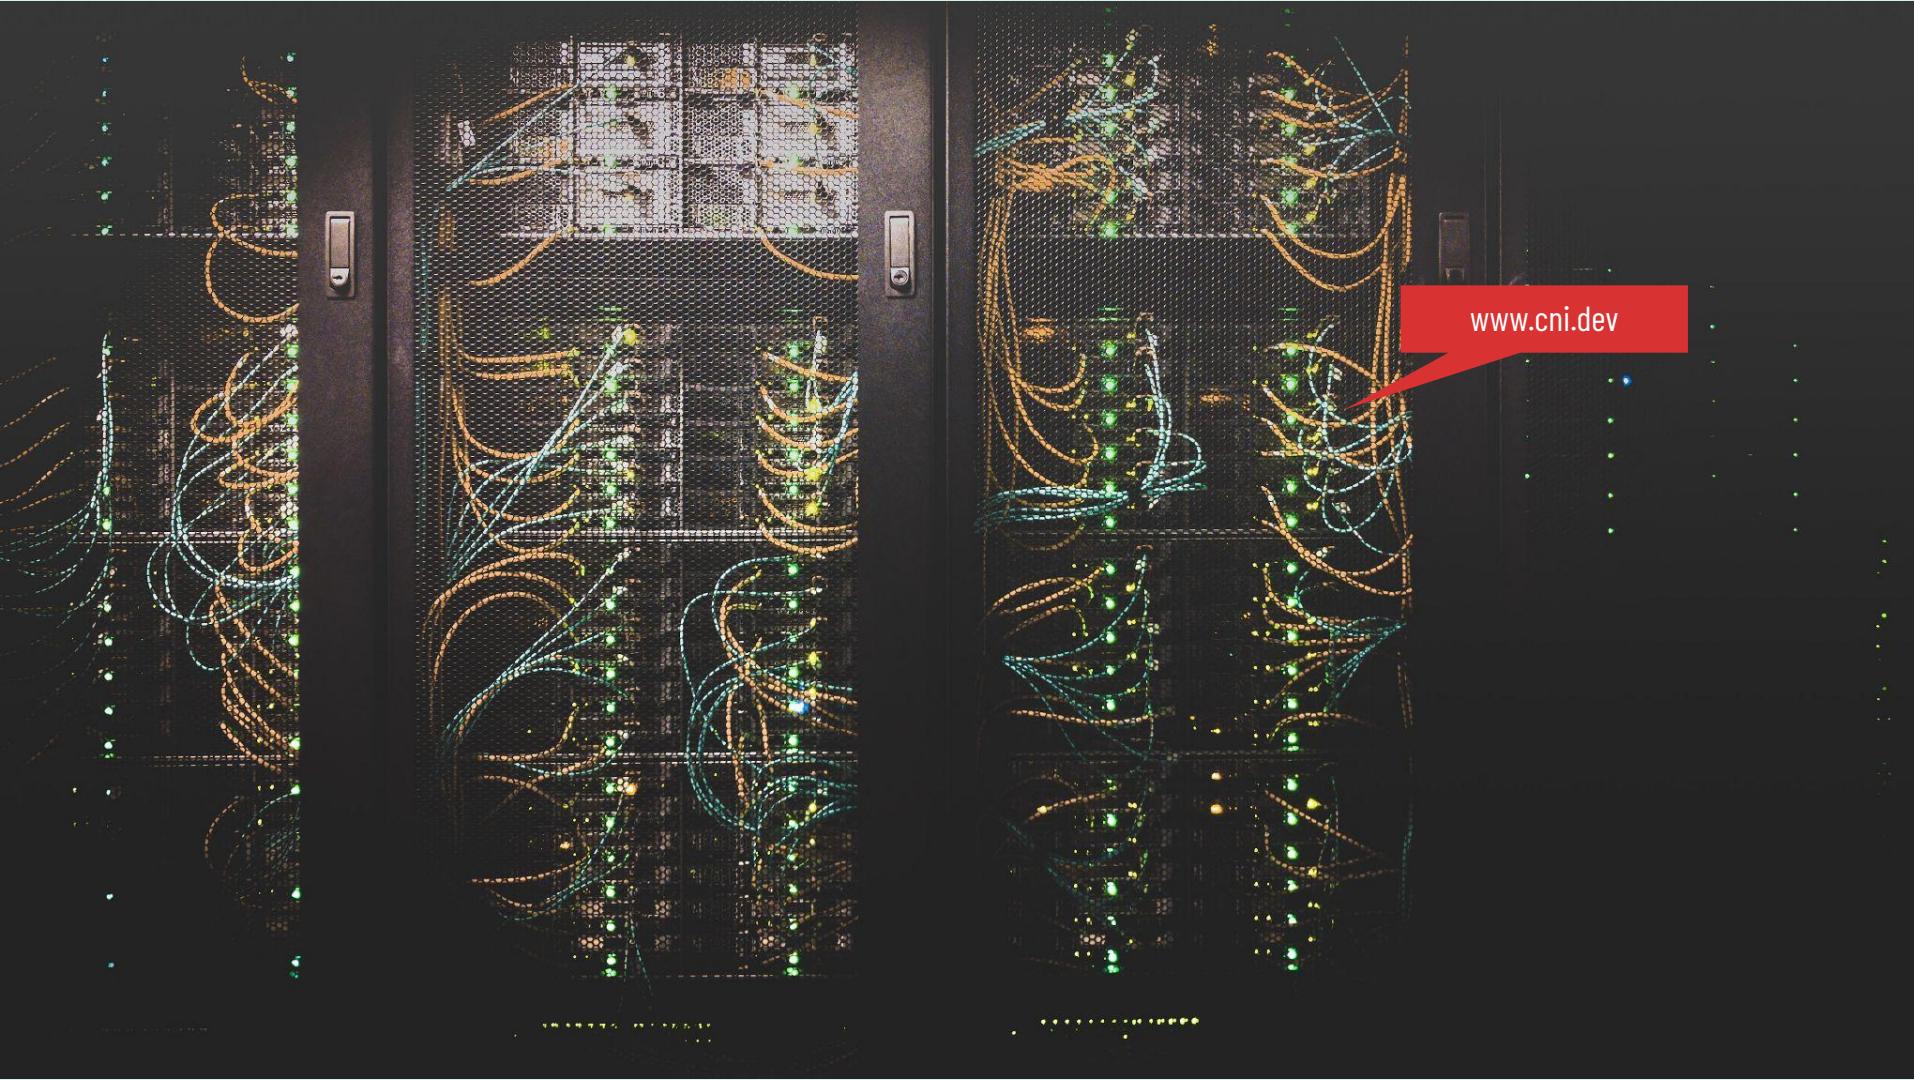
```
az aks create \
--resource-group demo-rg \
--name demo-cluster
```



# Network Plugin

```
az aks create \
--resource-group demo-rg \
--name demo-cluster \
--network-plugin {kubenet, azure}
```





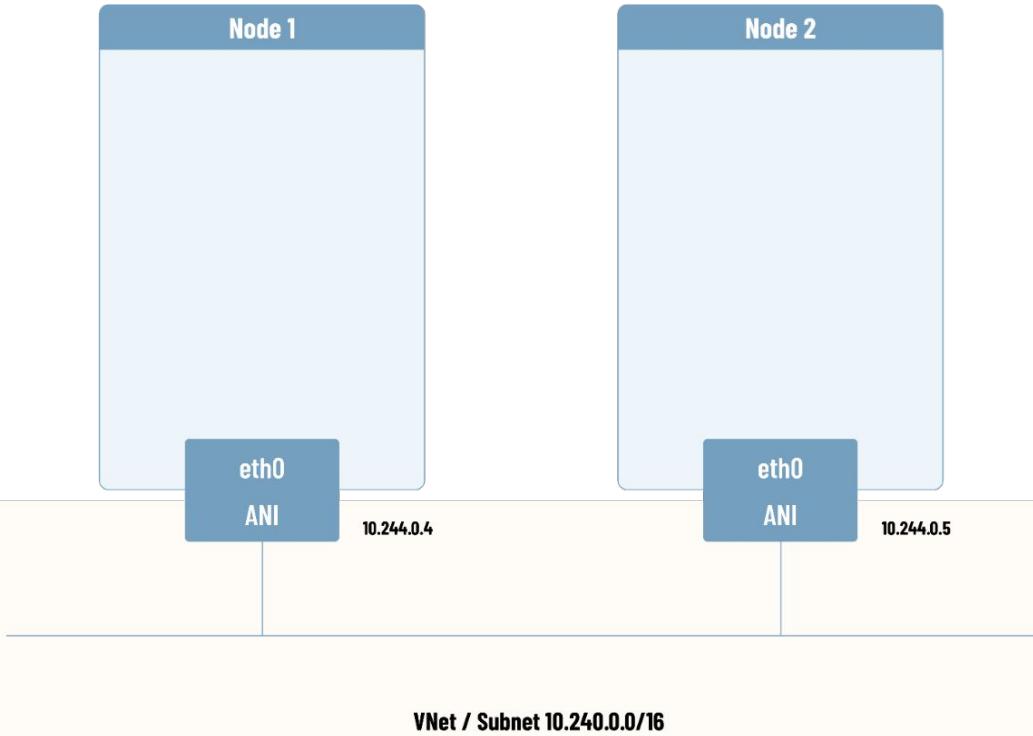
www.cni.dev



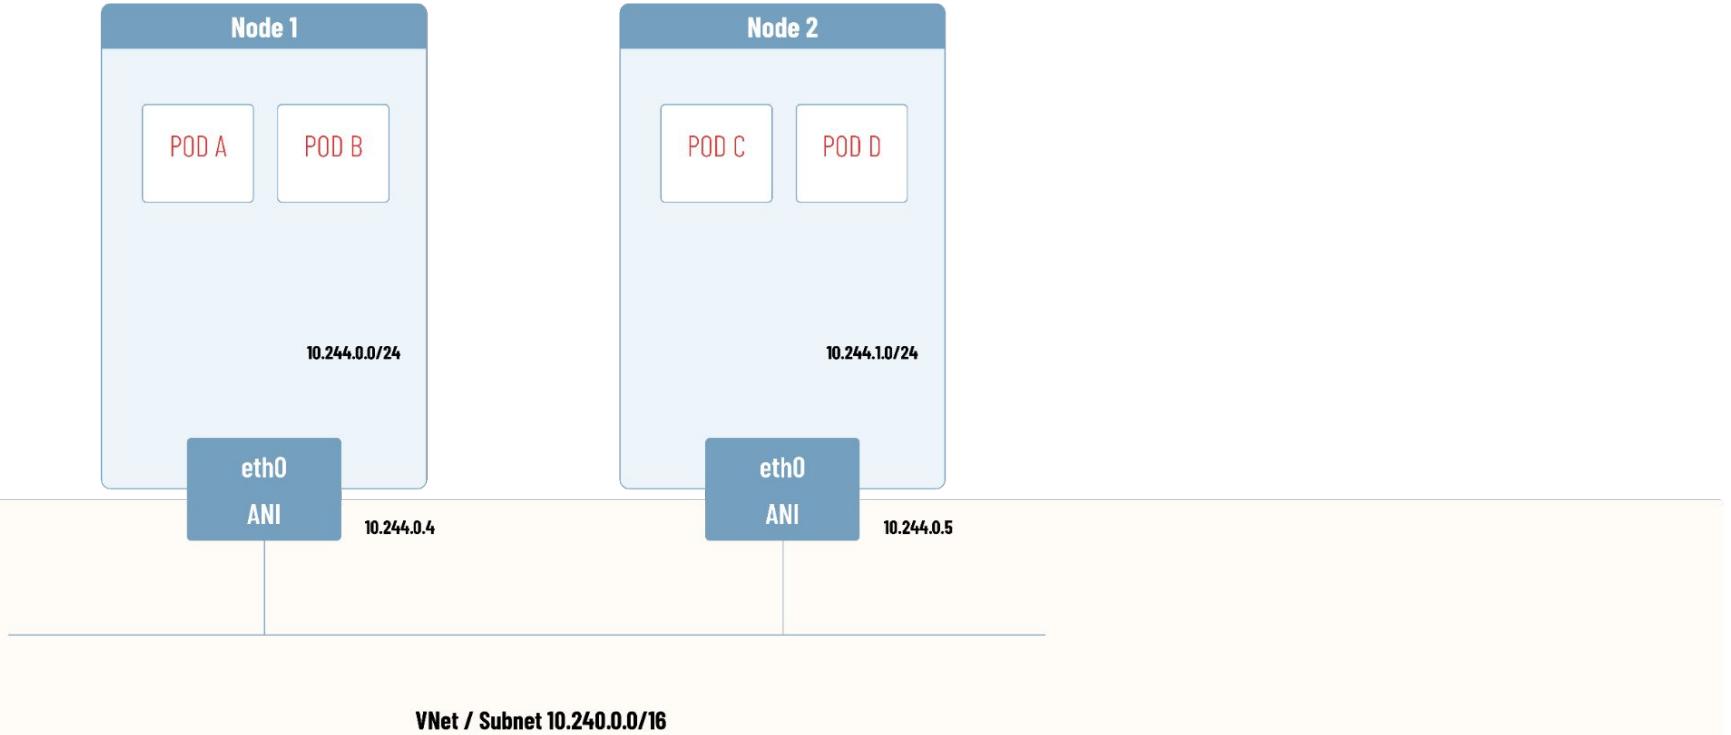
# Basic networking

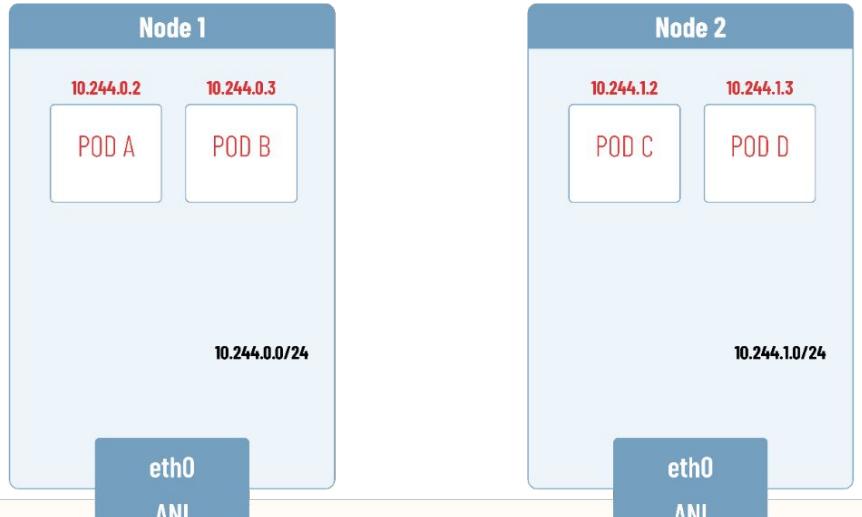
## kubenet

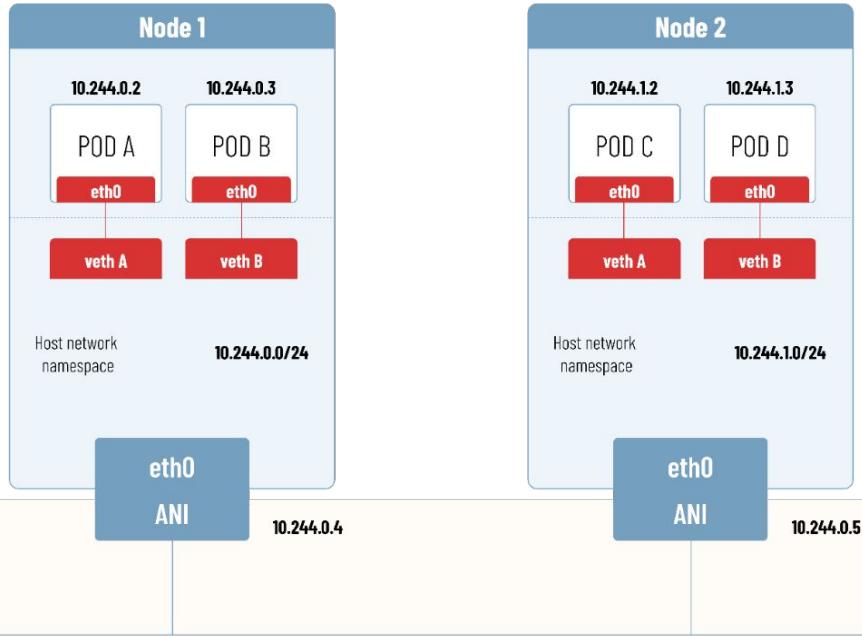


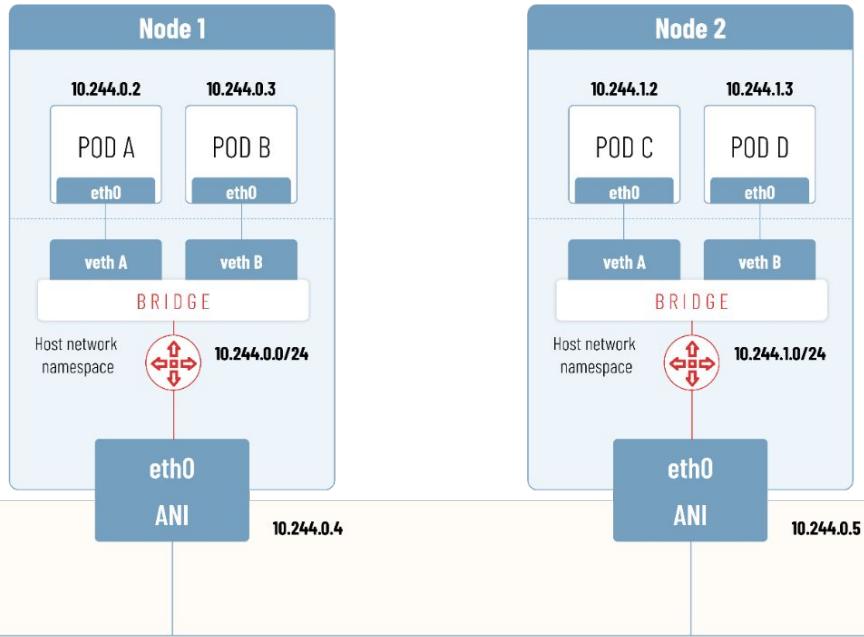










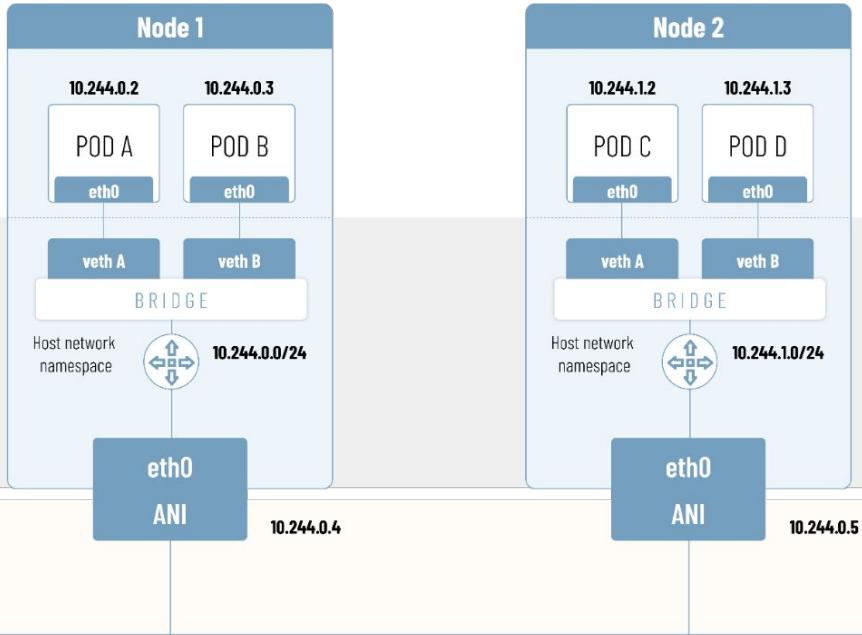


VNet / Subnet 10.240.0.0/16



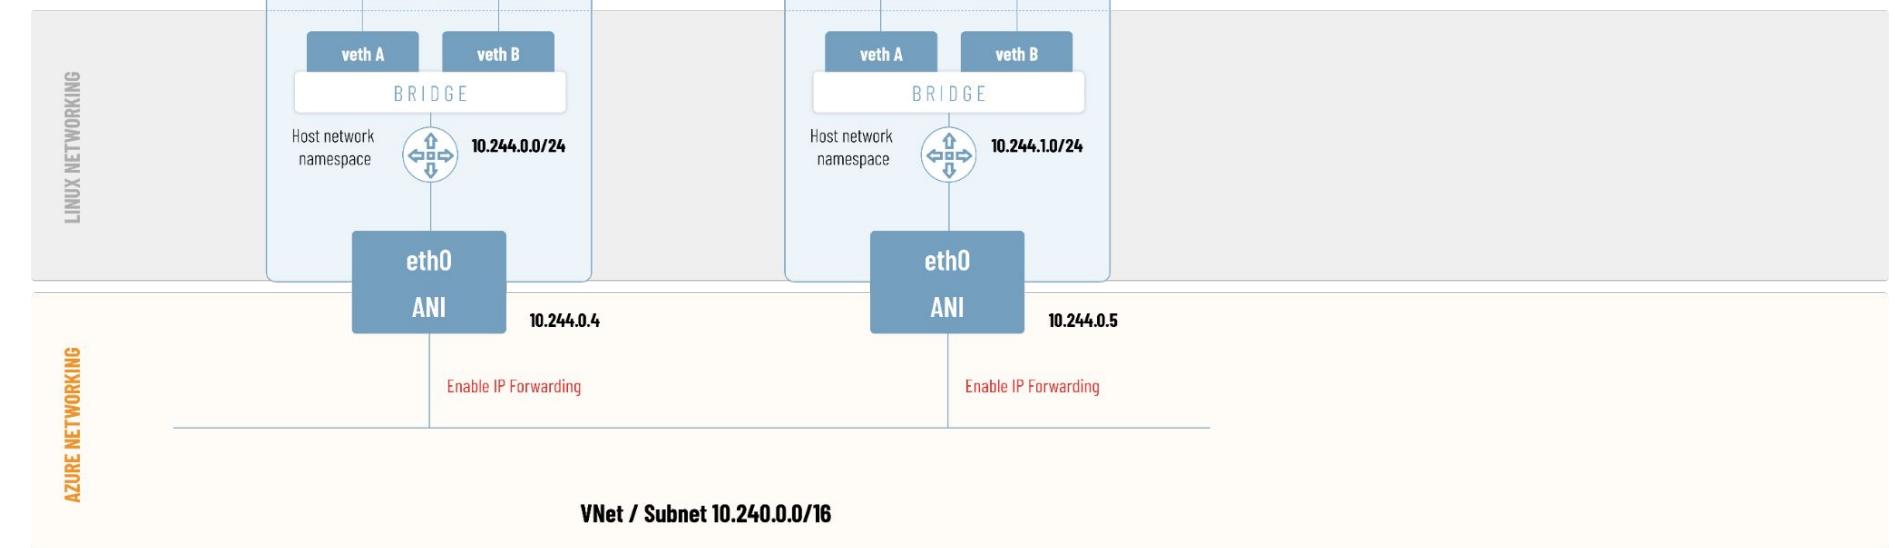
## LINUX NETWORKING

## AZURE NETWORKING

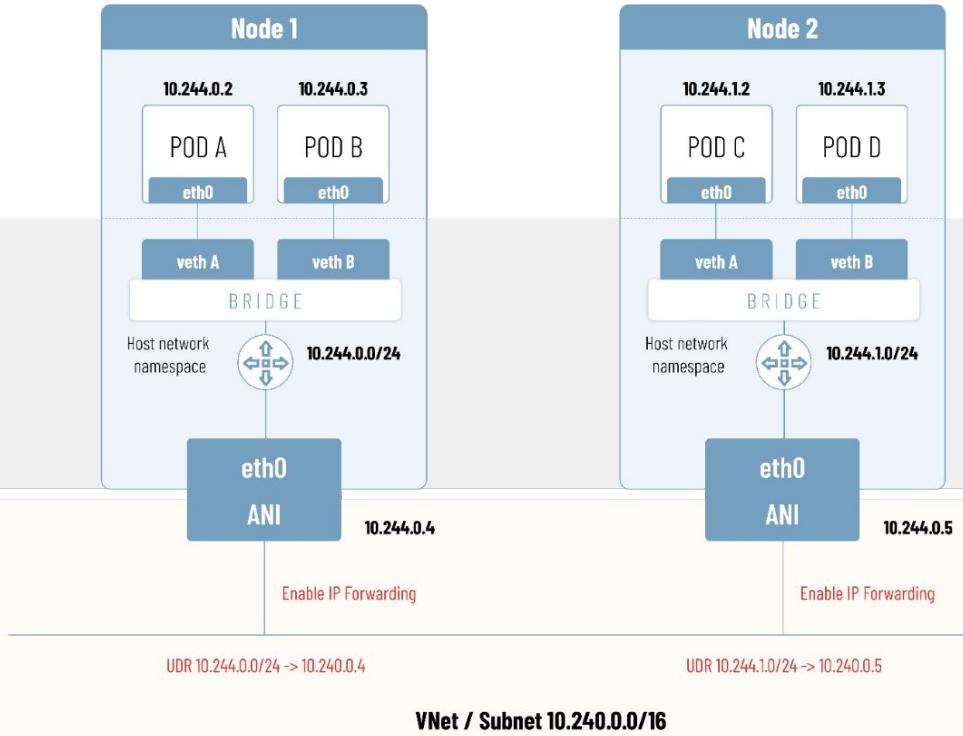


VNet / Subnet **10.240.0.0/16**



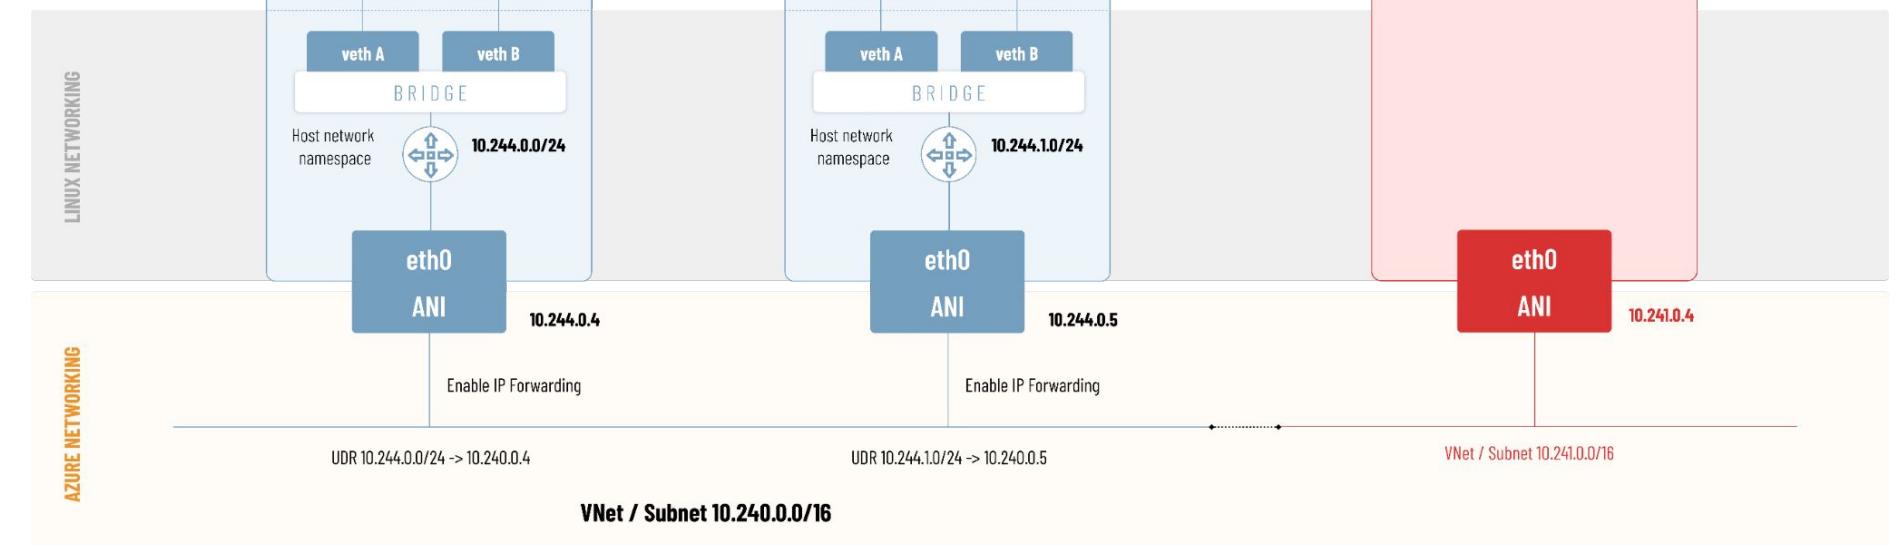


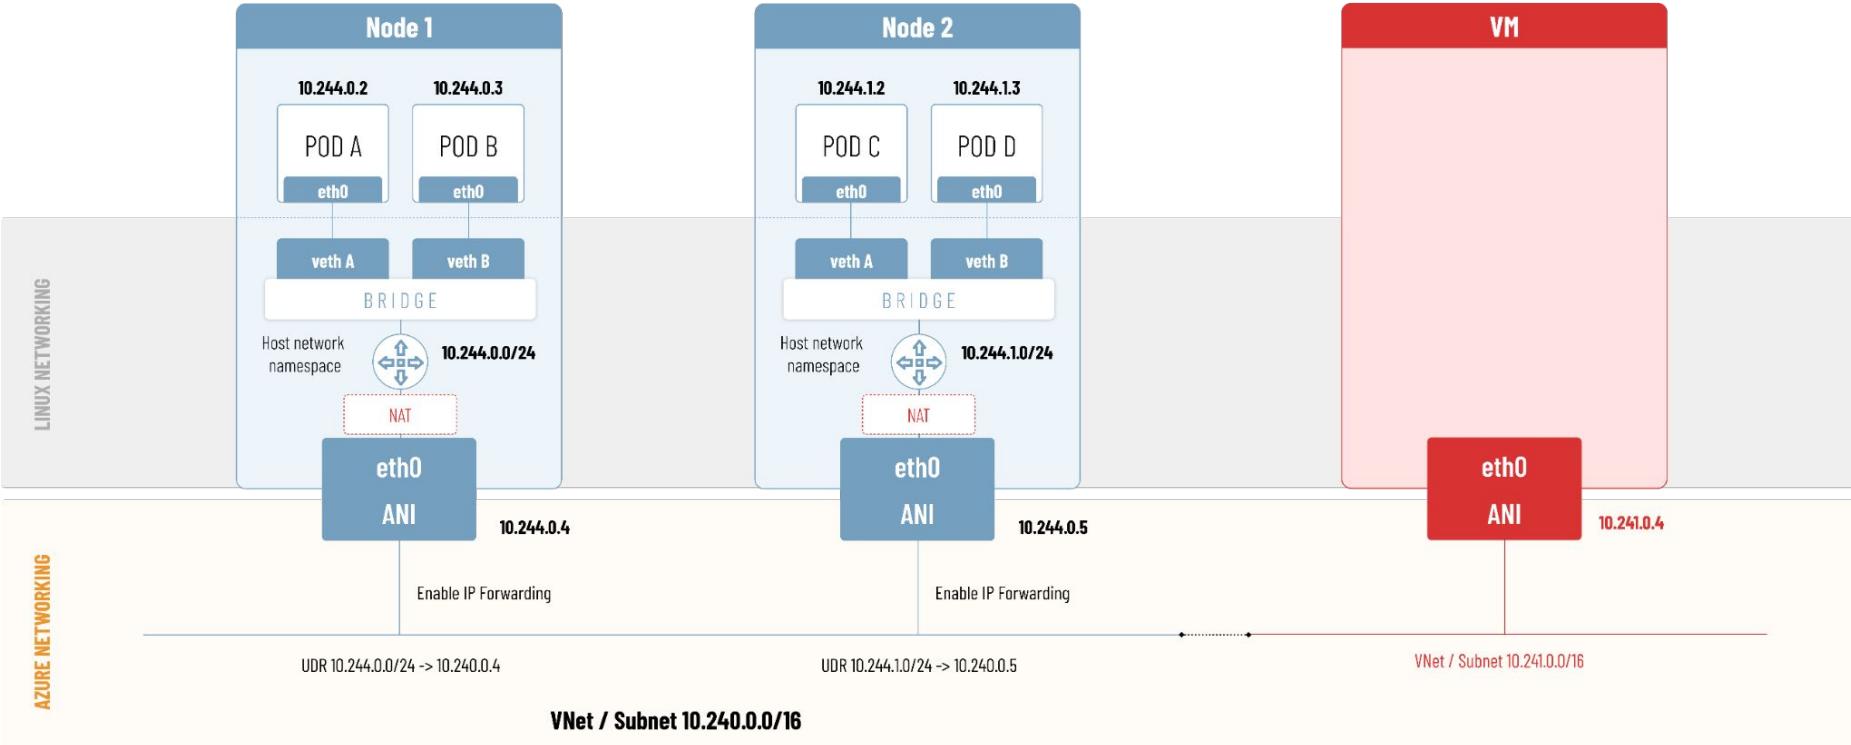
## LINUX NETWORKING



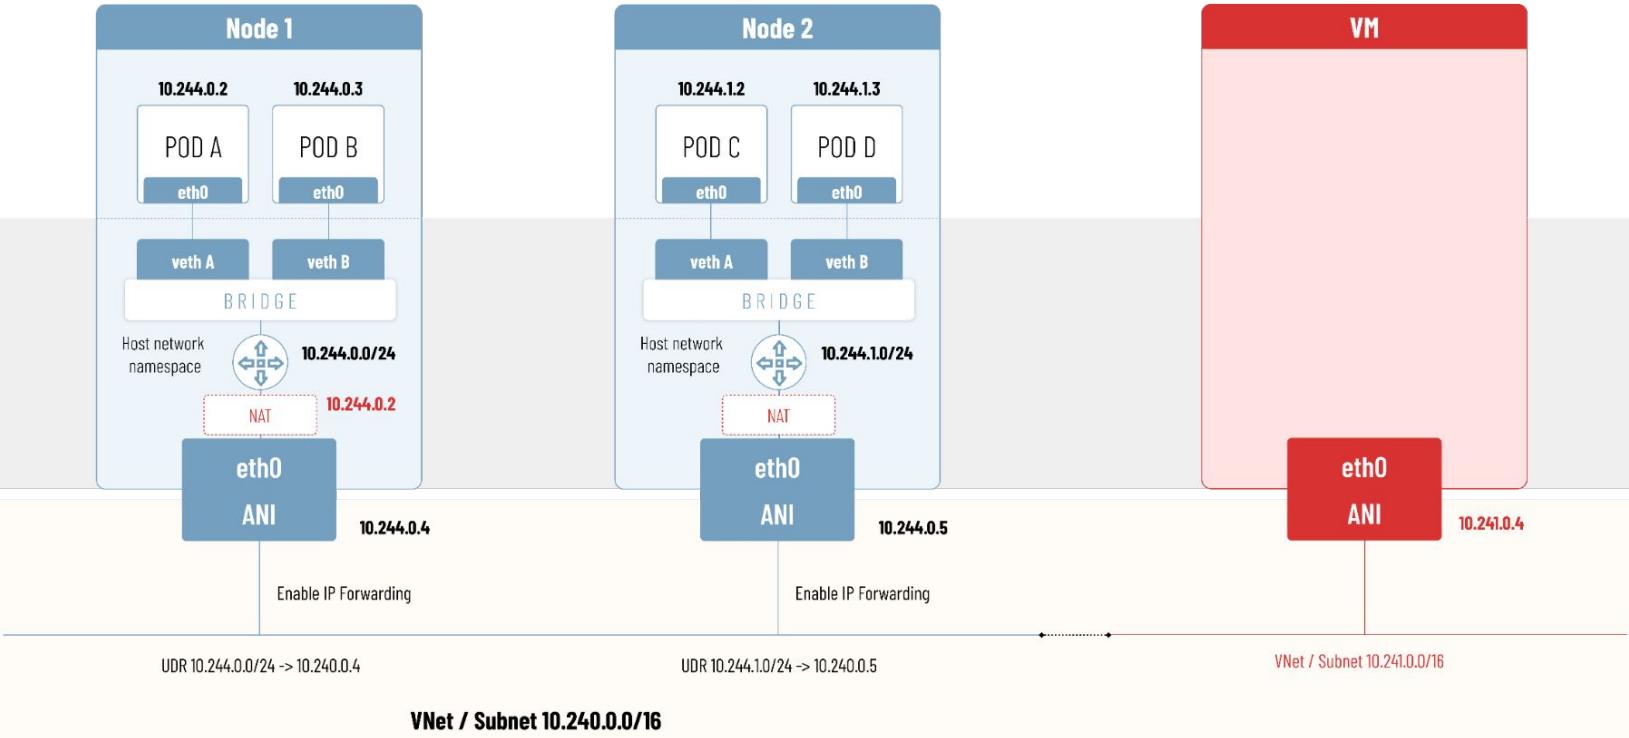
## AZURE NETWORKING







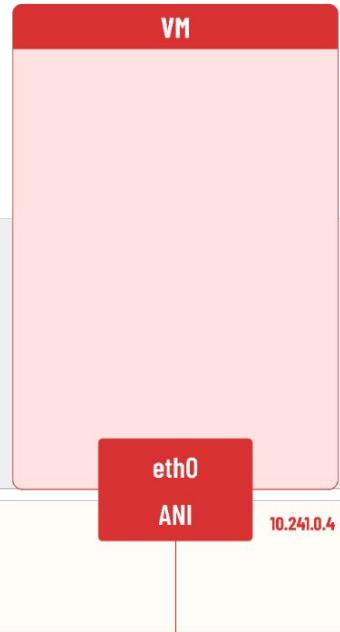
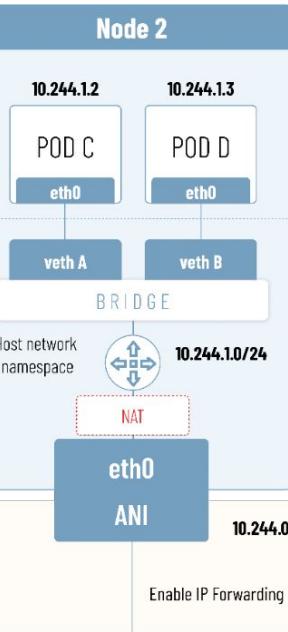
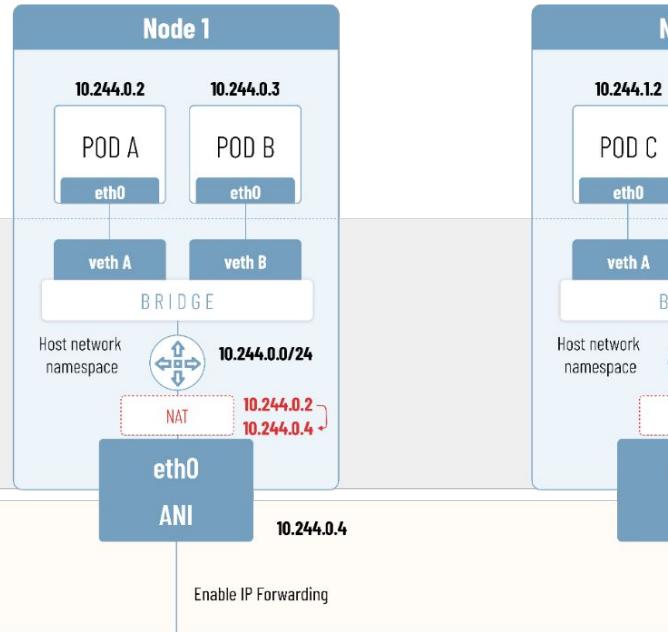
## LINUX NETWORKING



## AZURE NETWORKING



## LINUX NETWORKING



## AZURE NETWORKING

UDR 10.244.0.0/24 -> 10.240.0.4

UDR 10.244.1.0/24 -> 10.240.0.5

VNet / Subnet 10.241.0.0/16

**VNet / Subnet 10.240.0.0/16**

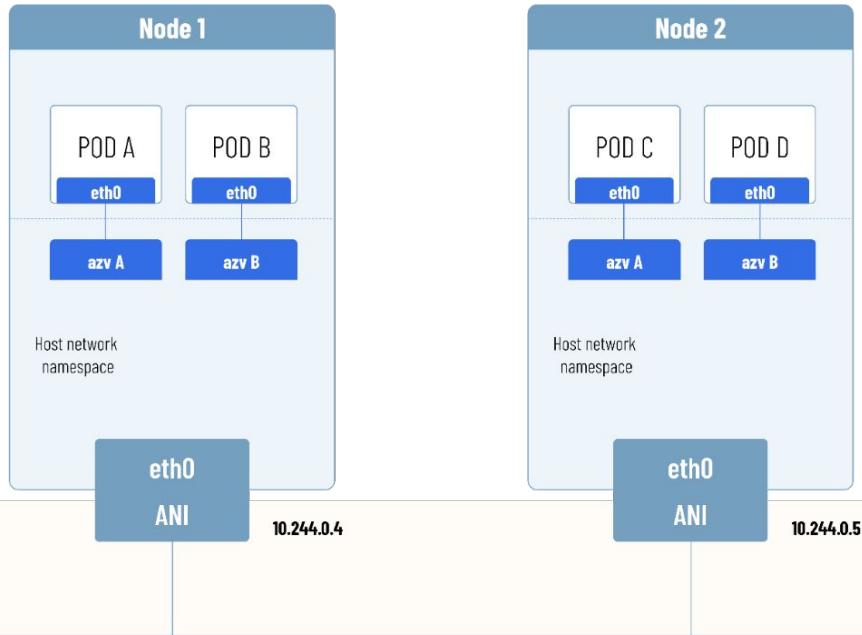


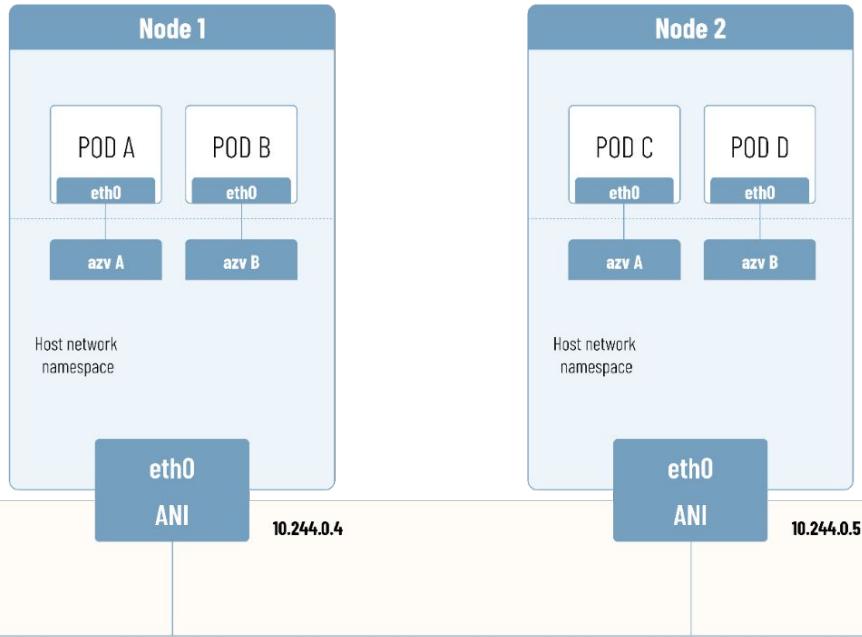


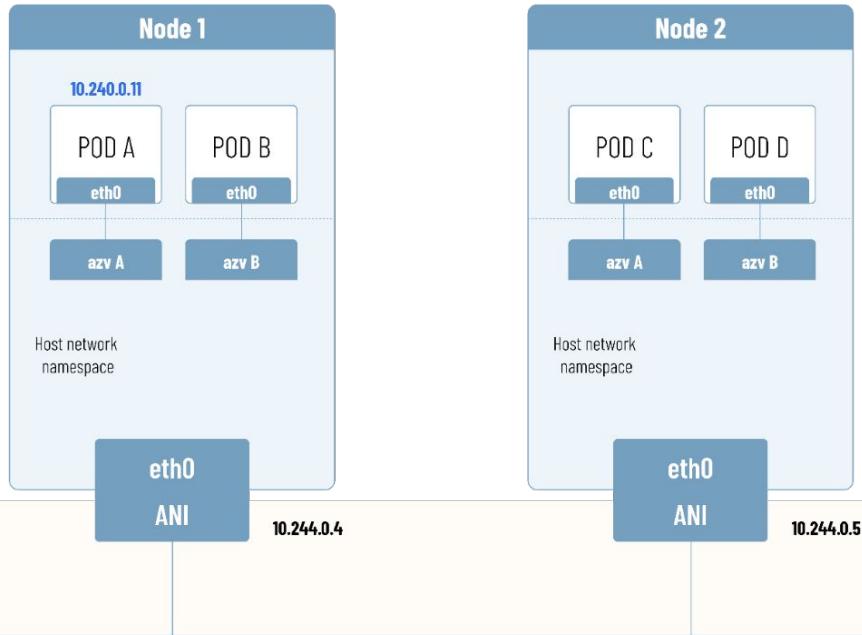
# Advanced networking

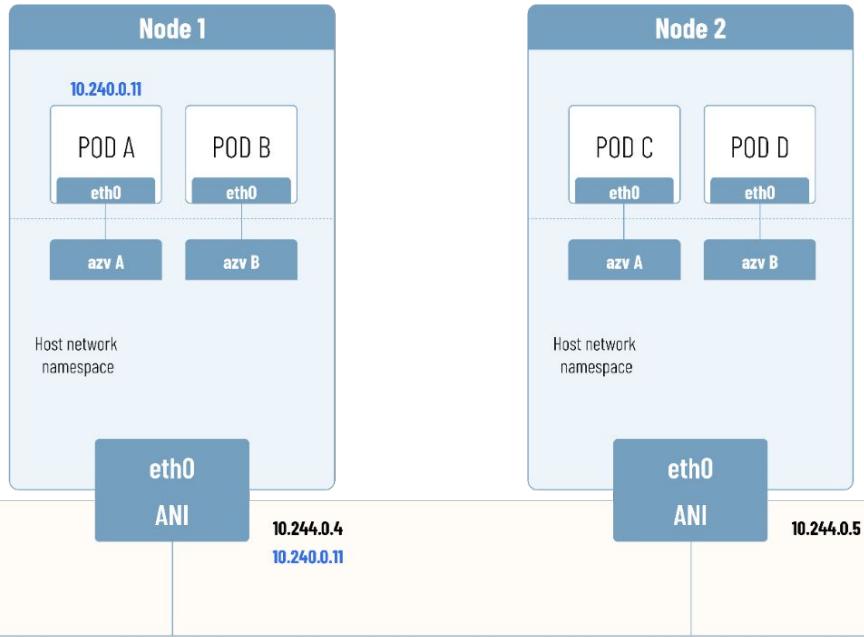
## Azure CNI





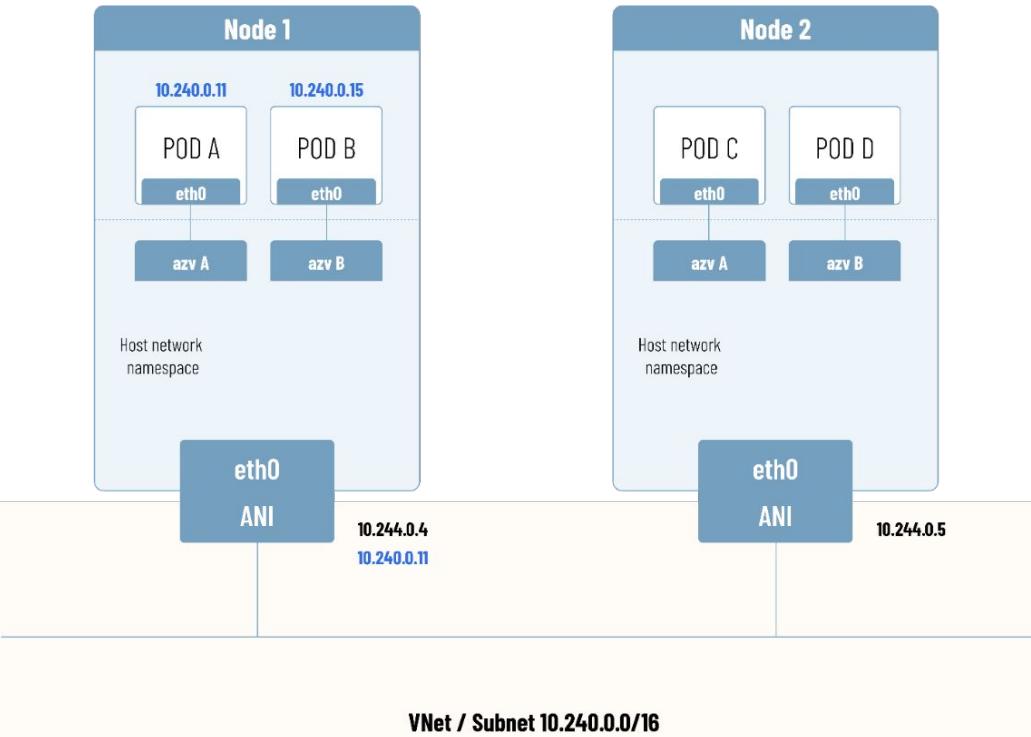


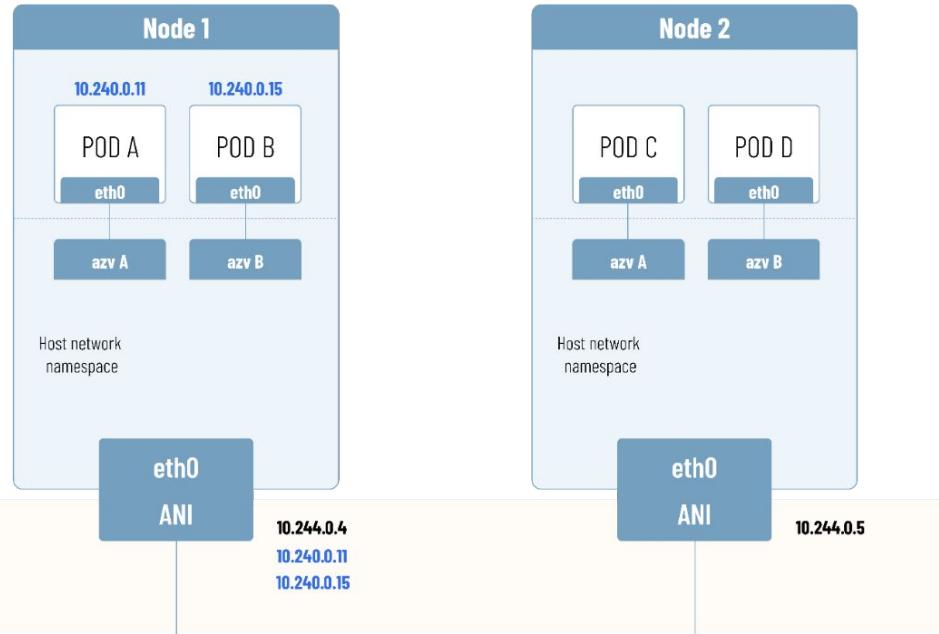


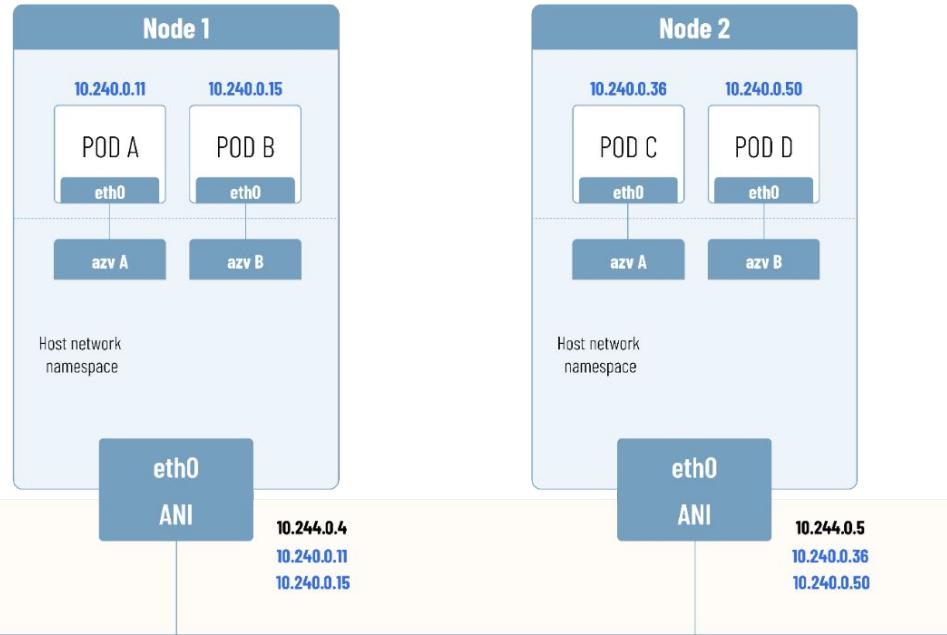


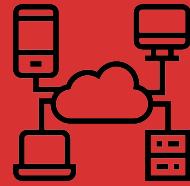
VNet / Subnet **10.240.0.0/16**











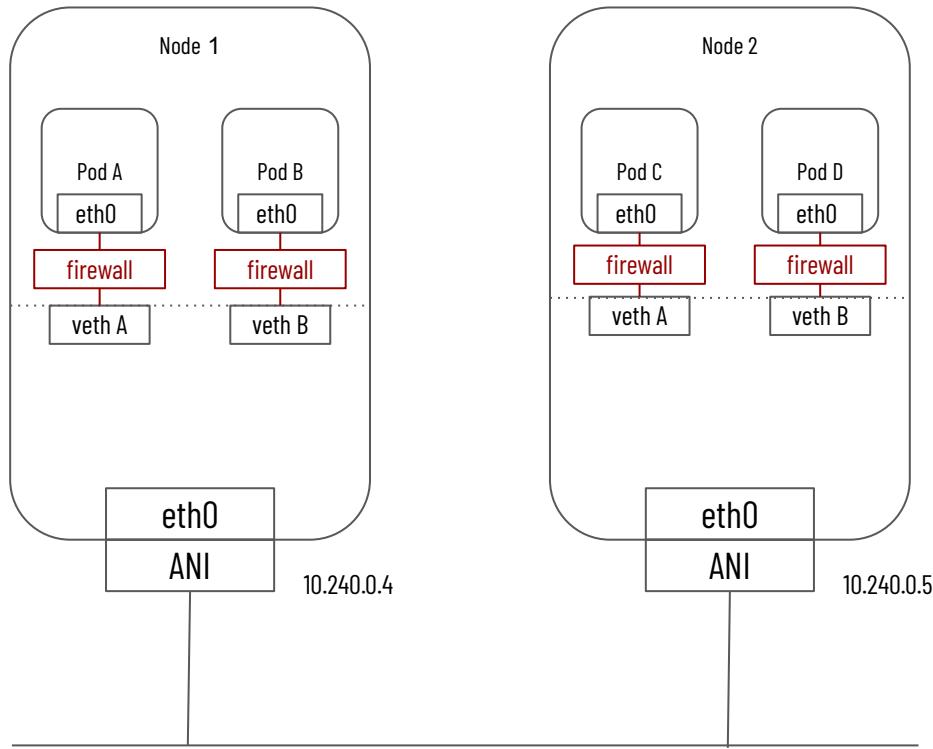
# Network policies



# Network Policy

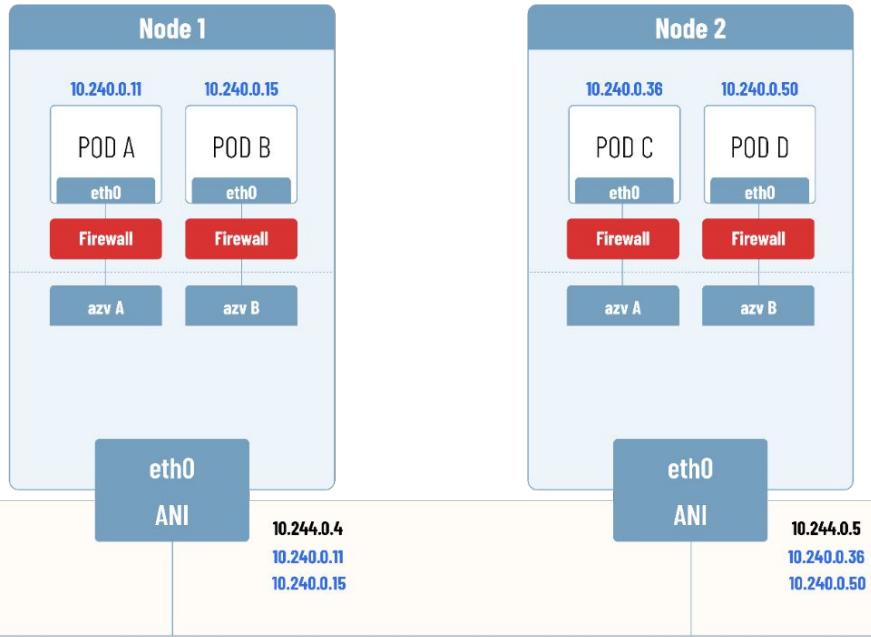
```
az aks create \
--resource-group demo-rg \
--name demo-cluster \
--network-plugin {kubenet, azure} \
--network-policy {azure, calico}
```





VNet / Subnet 10.240.0.0/16



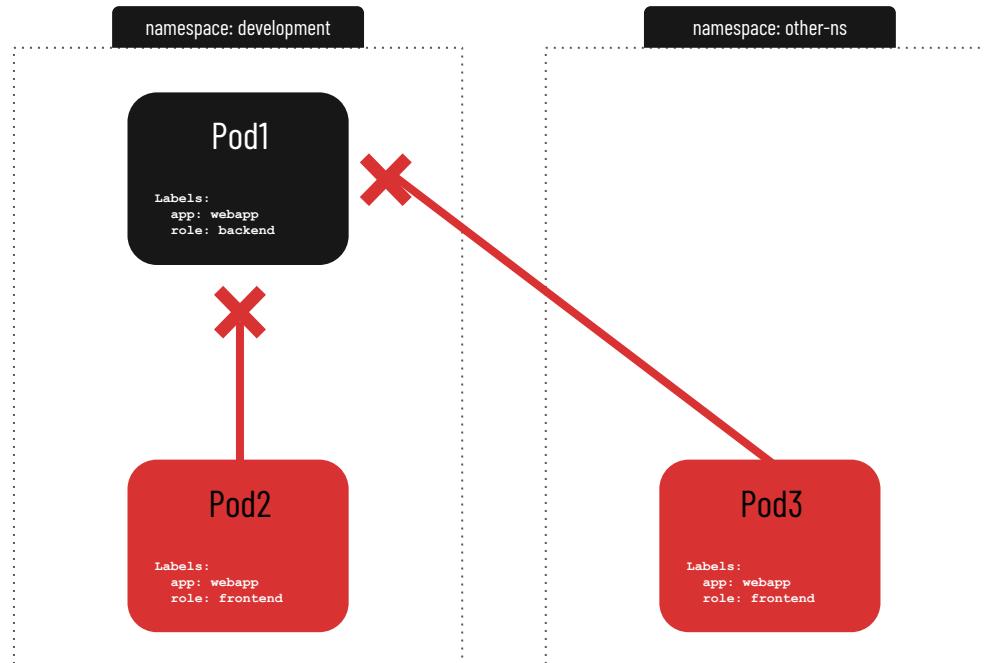


VNet / Subnet 10.240.0.0/16



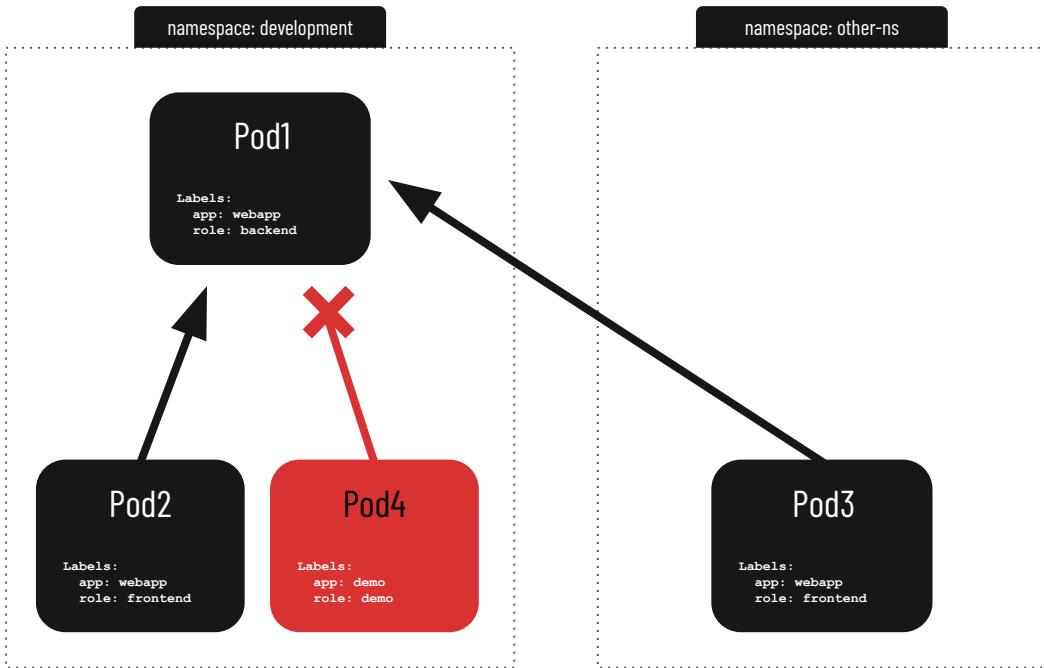
# Deny all inbound traffic to a pod

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: backend-policy
  namespace: development
spec:
  podSelector:
    matchLabels:
      app: webapp
      role: backend
  ingress: []
```



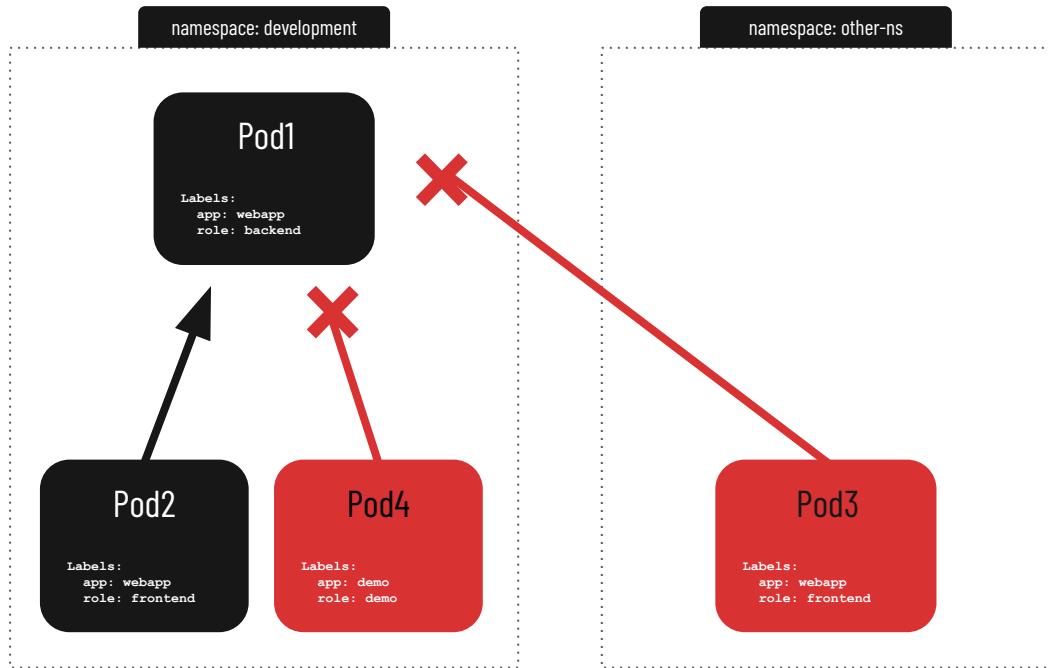
# Allow inbound traffic based on a pod label

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: backend-policy
  namespace: development
spec:
  podSelector:
    matchLabels:
      app: webapp
      role: backend
  ingress:
  - from:
    - namespaceSelector: {}
      podSelector:
        matchLabels:
          app: webapp
          role: frontend
```



# Allow traffic only from within a defined namespace

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: backend-policy
  namespace: development
spec:
  podSelector:
    matchLabels:
      app: webapp
      role: backend
  ingress:
  - from:
    - namespaceSelector:
        matchLabels:
          purpose: development
  podSelector:
    matchLabels:
      app: webapp
      role: frontend
```



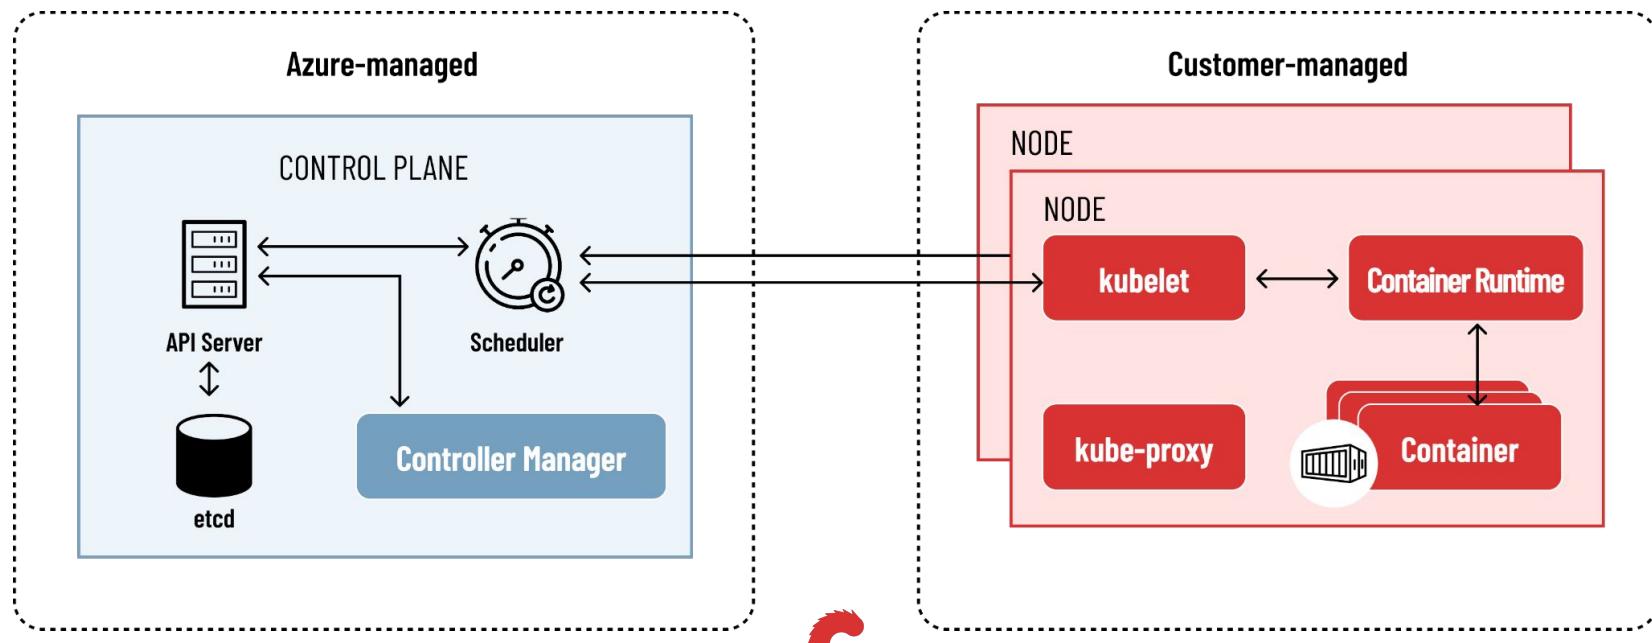


# Business continuity



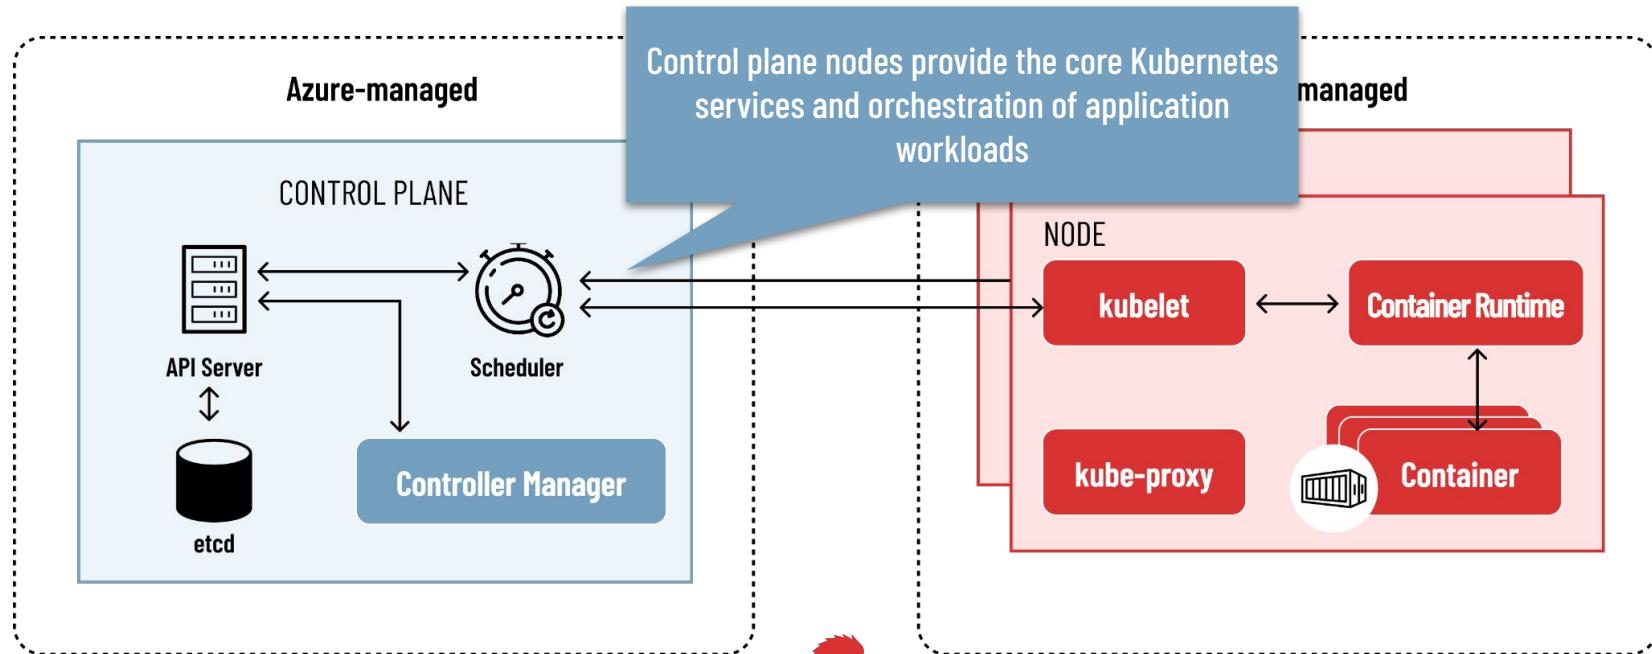
# Kubernetes cluster architecture

A Kubernetes cluster is divided into two components:



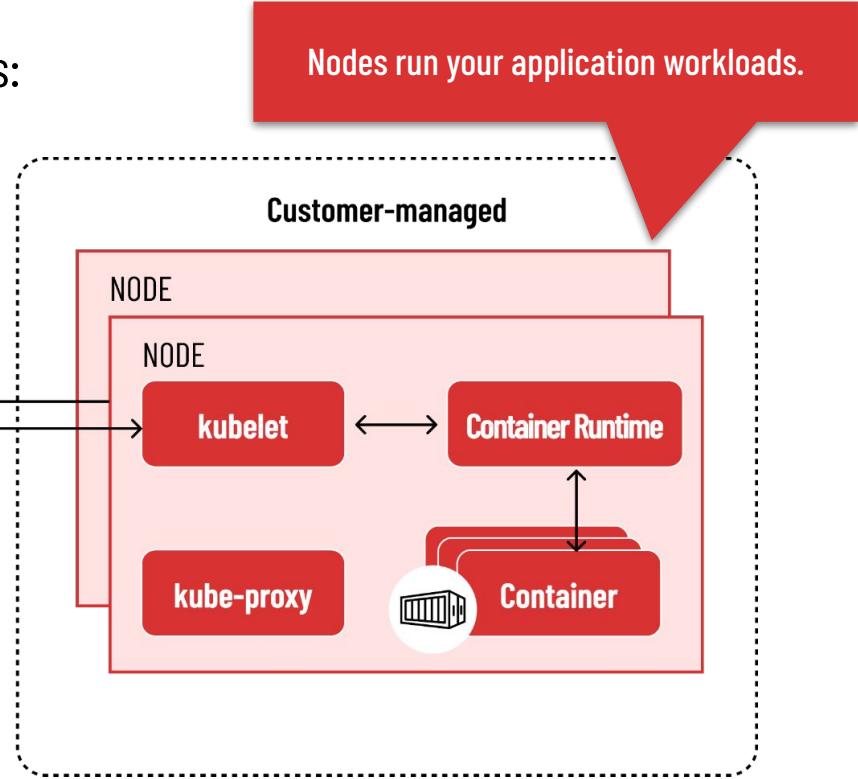
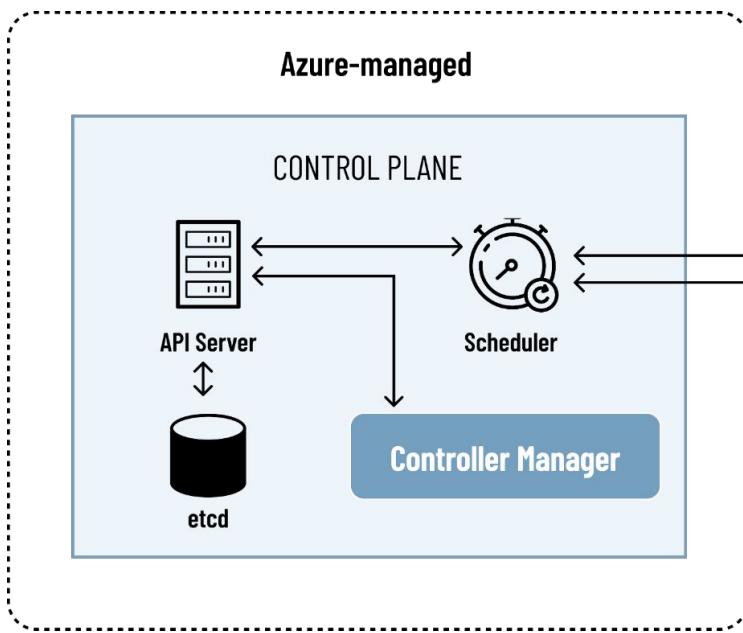
# Kubernetes cluster architecture

A Kubernetes cluster is divided into two components:

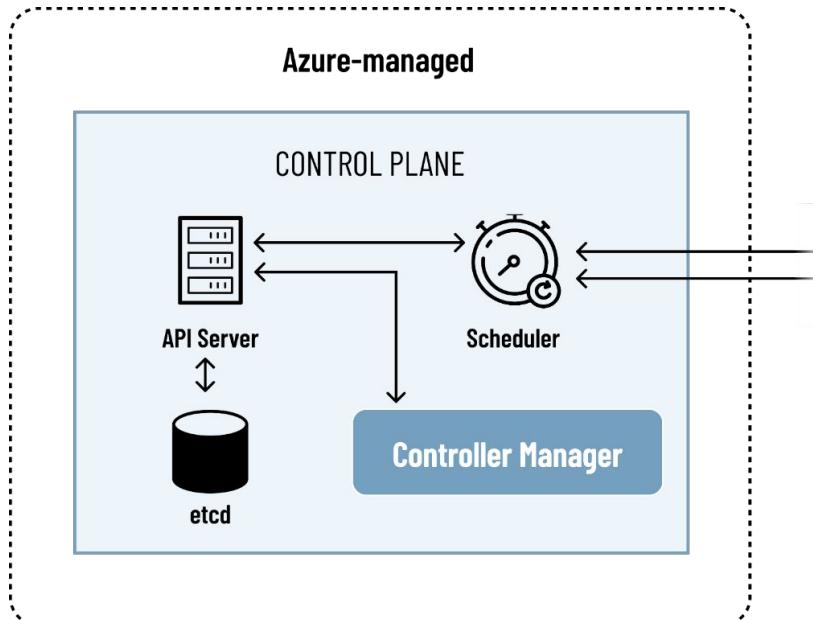


# Kubernetes cluster architecture

A Kubernetes cluster is divided into two components:



# Control plane



When you create an AKS cluster, a control plane is automatically created and configured.

This control plane is provided as a **managed Azure resource** abstracted from the user.

There's **no cost** for the control plane, only the nodes that are part of the AKS cluster.



# API server business SLA

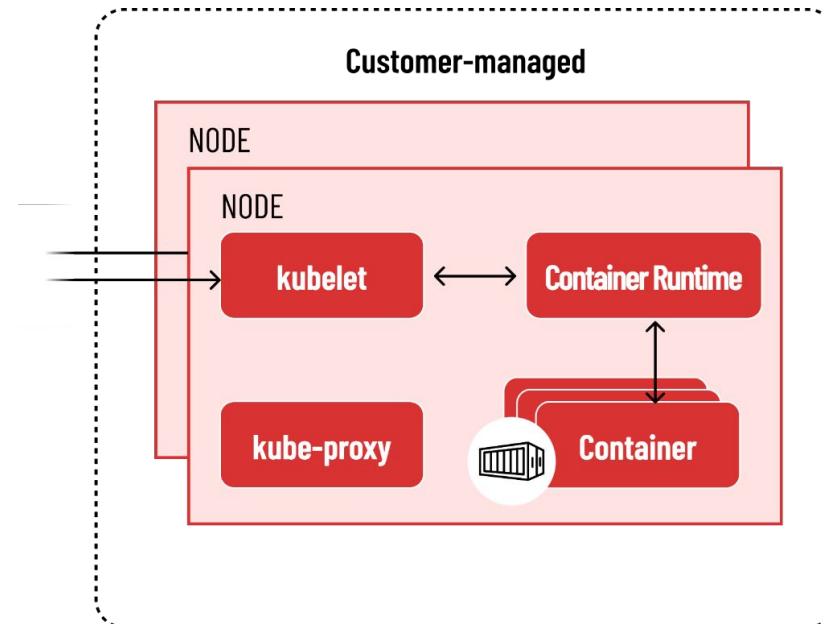
```
az aks create \
--resource-group demo-rg \
--name demo-cluster \
--uptime-sla
```



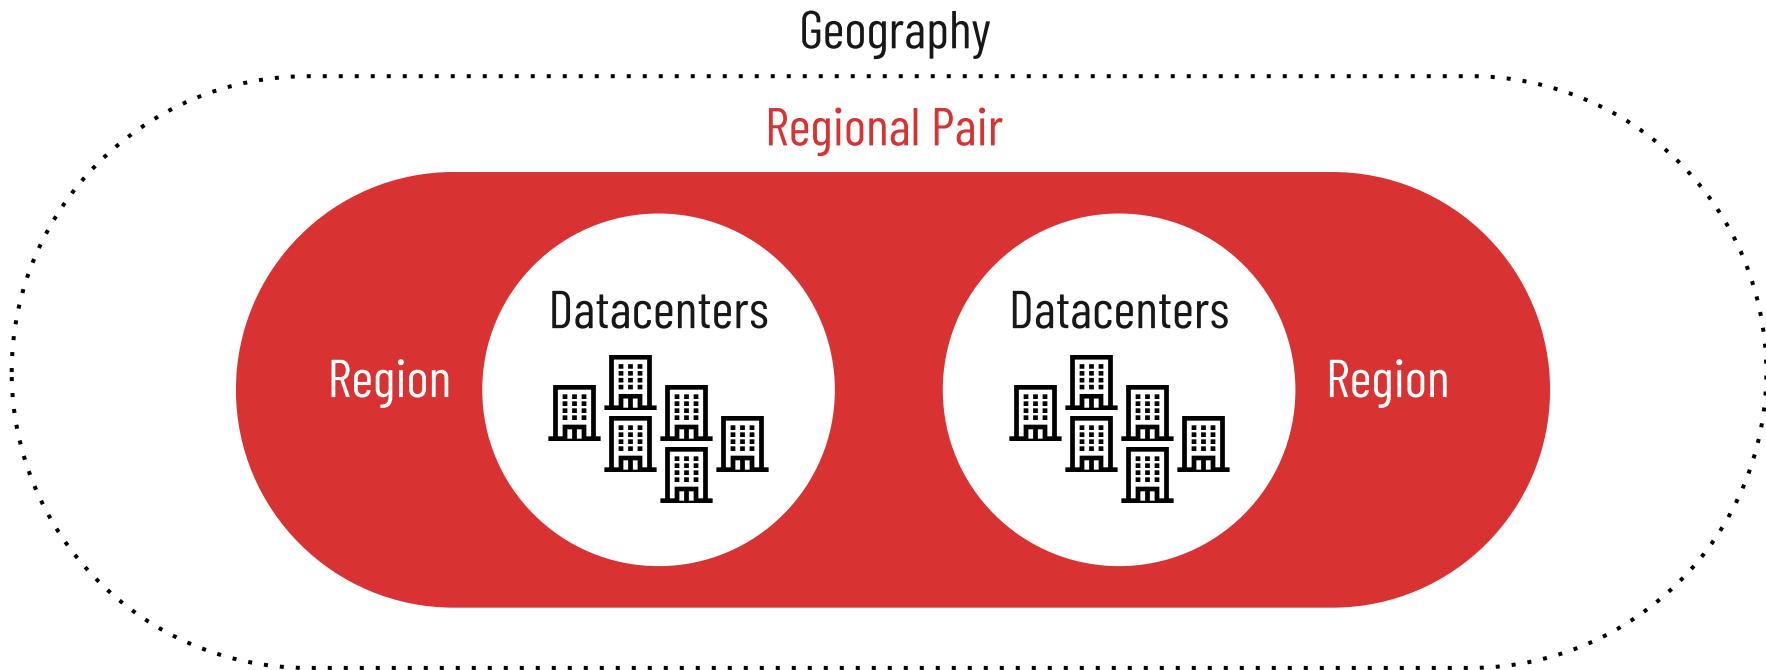
# Nodes and node pools

To run your applications and supporting services, you need a Kubernetes node.

An AKS cluster has one or more nodes, which is an Azure virtual machine (VM) that runs the [Kubernetes node components](#) and container runtime



# Allow traffic only from within a defined namespace

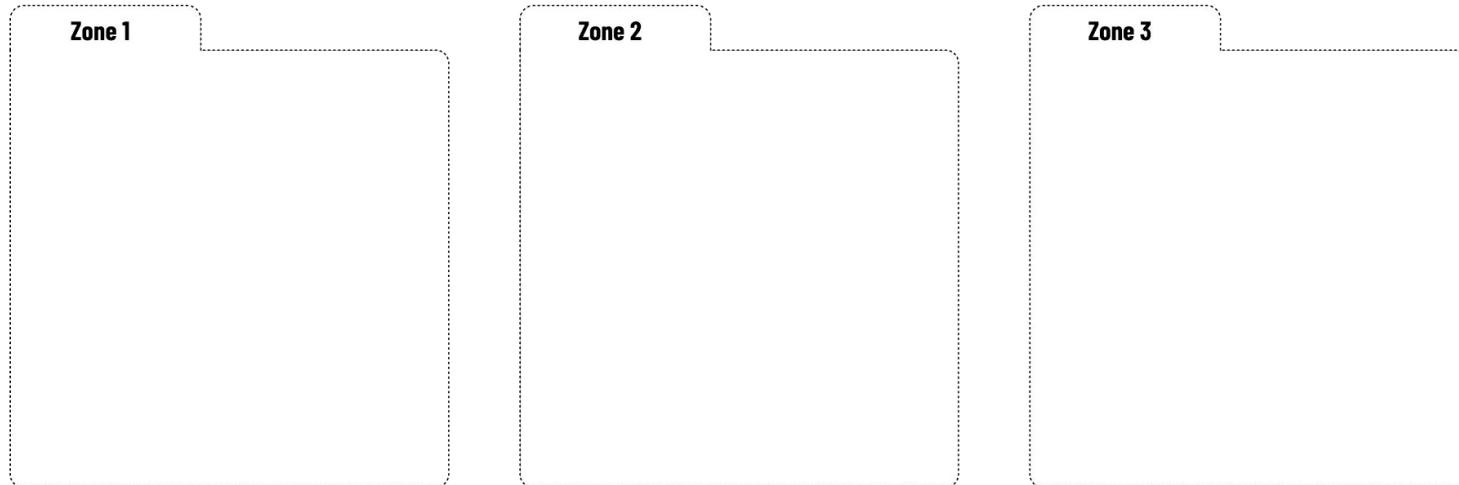


# Availability zones

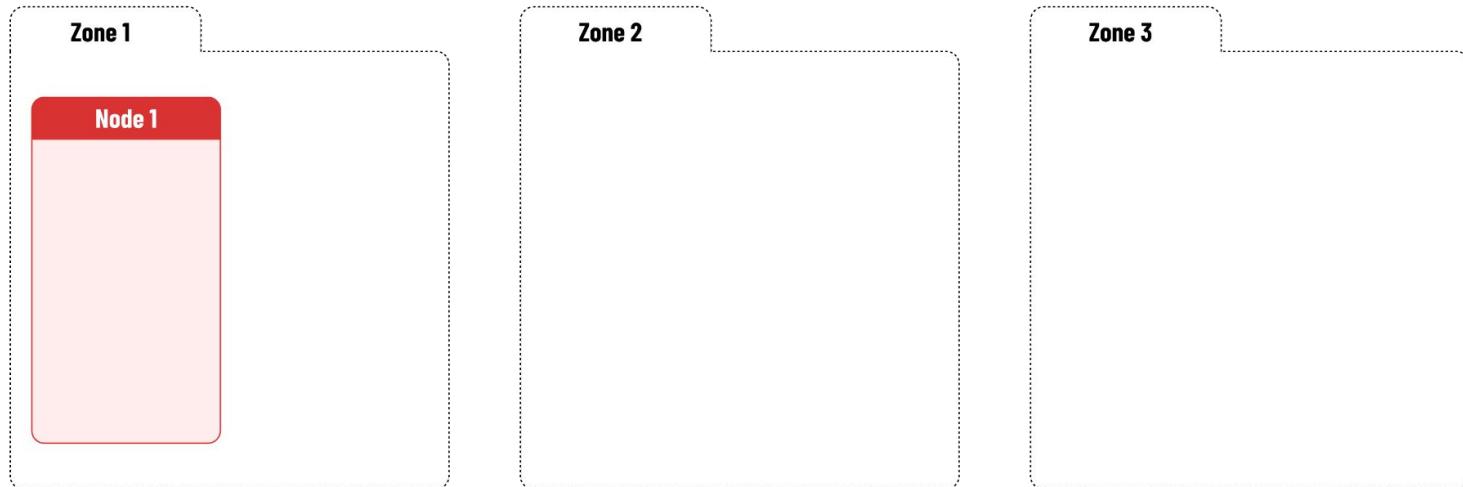
```
az aks create \
--resource-group demo-rg \
--name demo-cluster \
--zones {1, 2, 3}
```



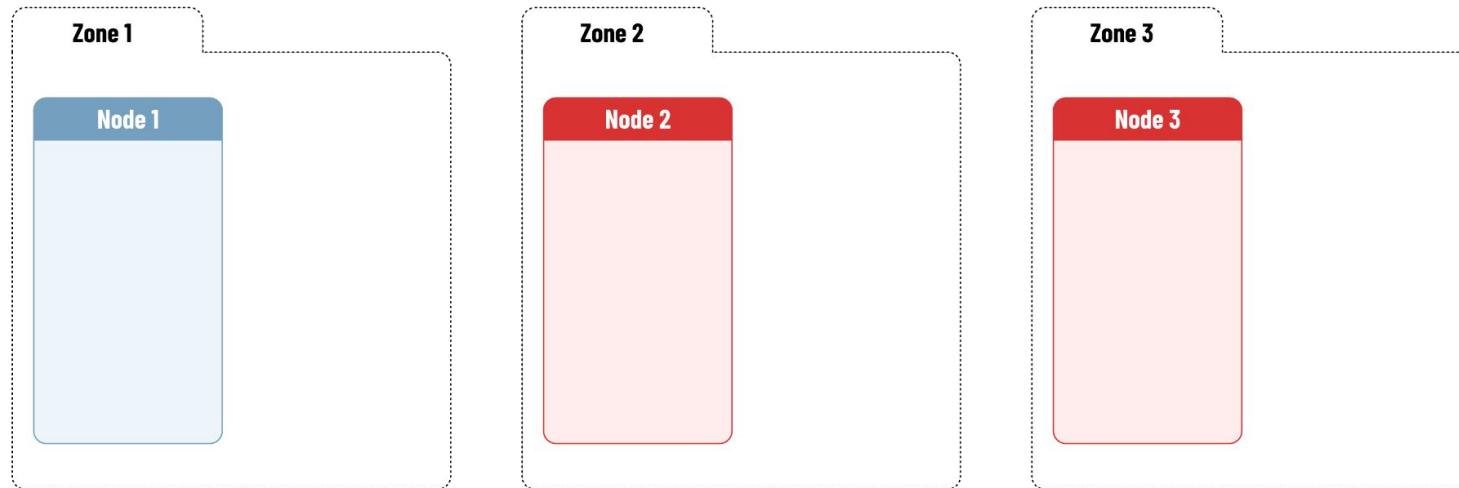
# AKS and Availability Zones



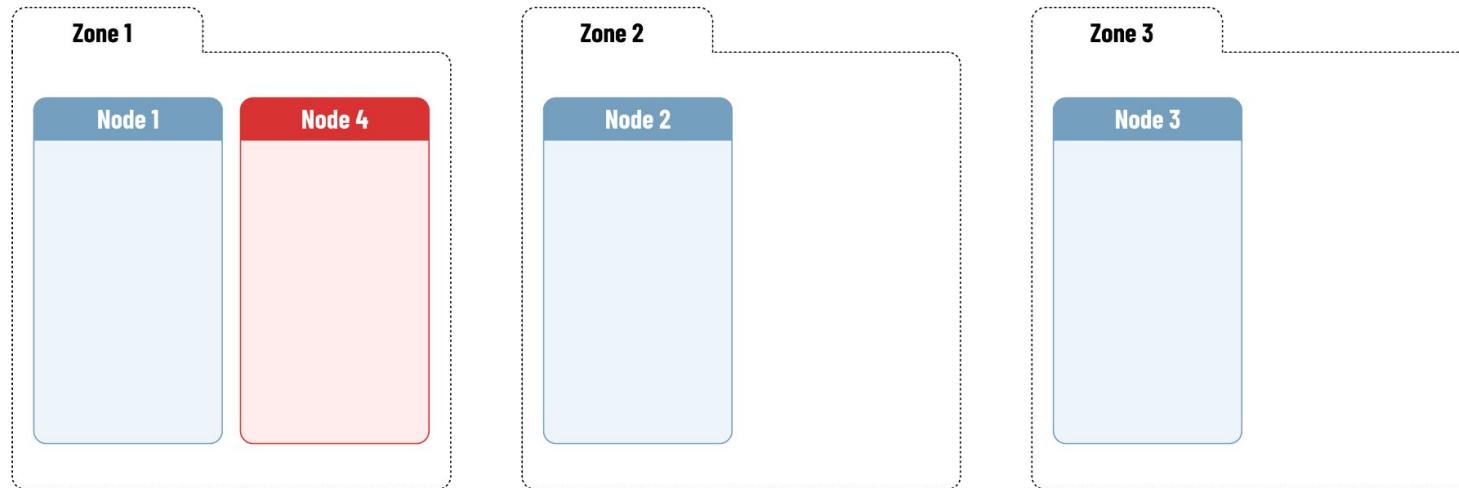
# AKS and Availability Zones



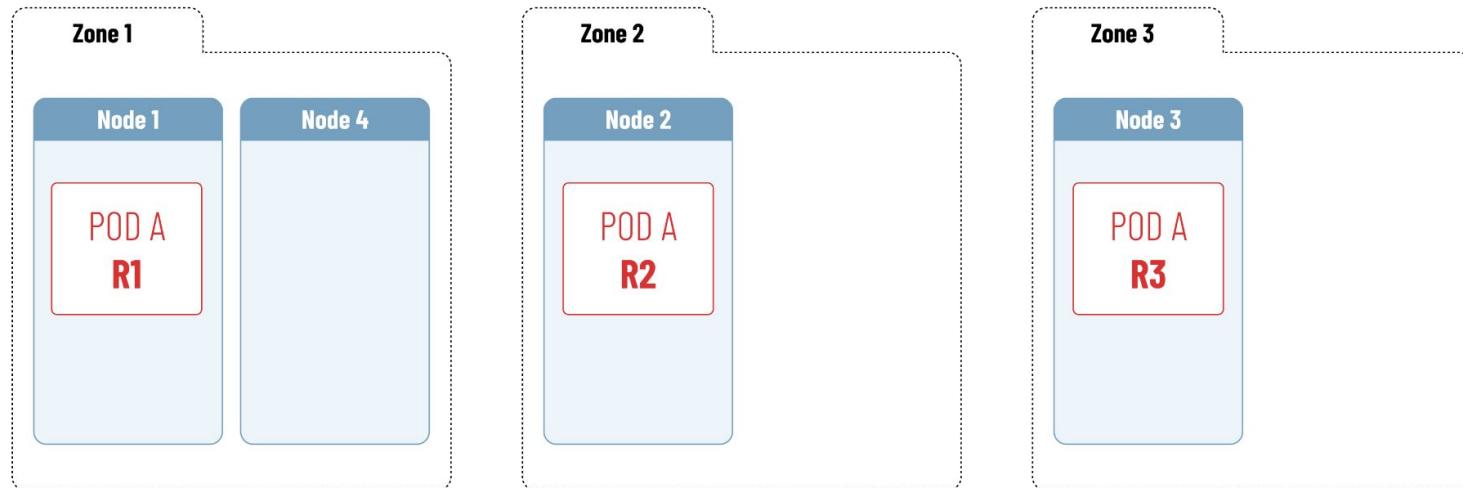
# AKS and Availability Zones



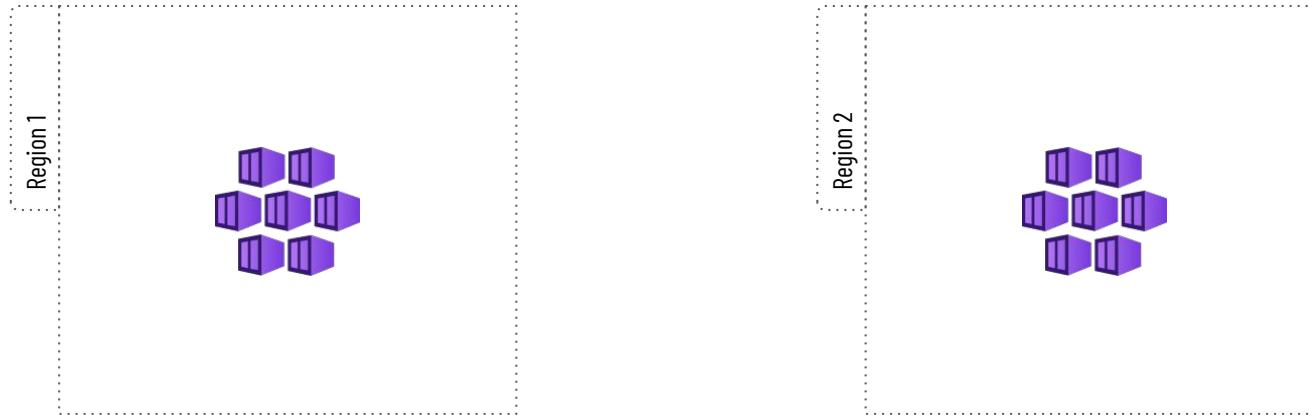
# AKS and Availability Zones



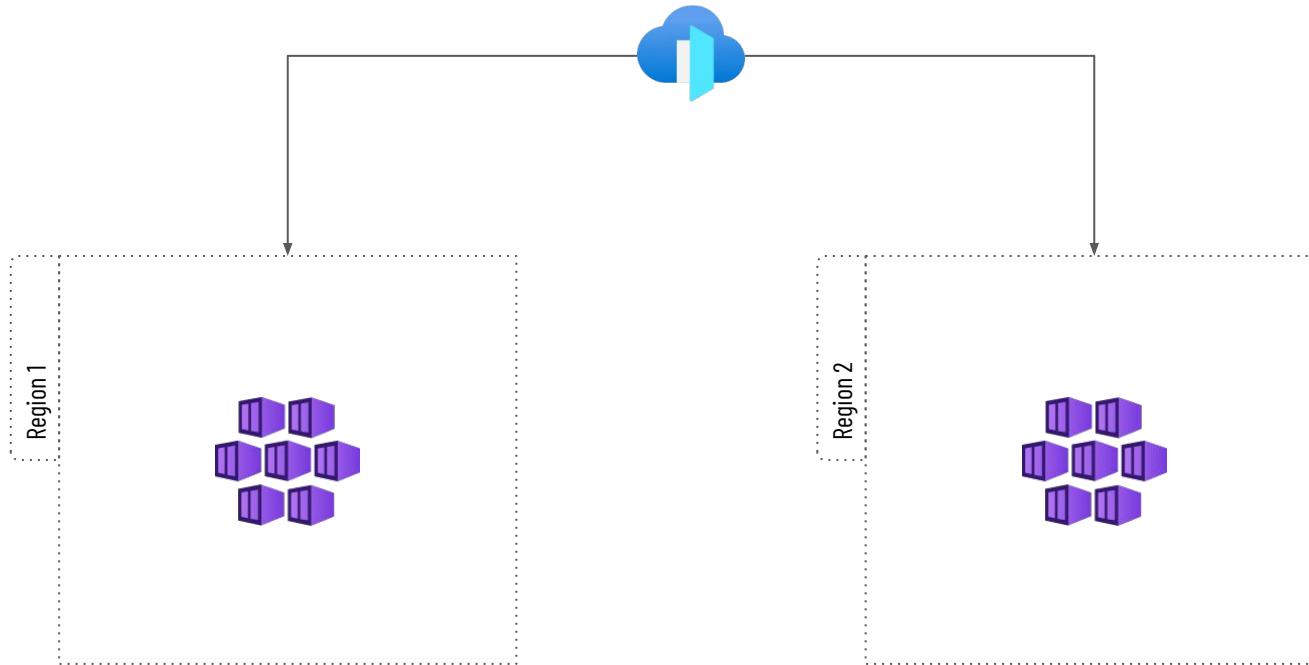
# AKS and Availability Zones



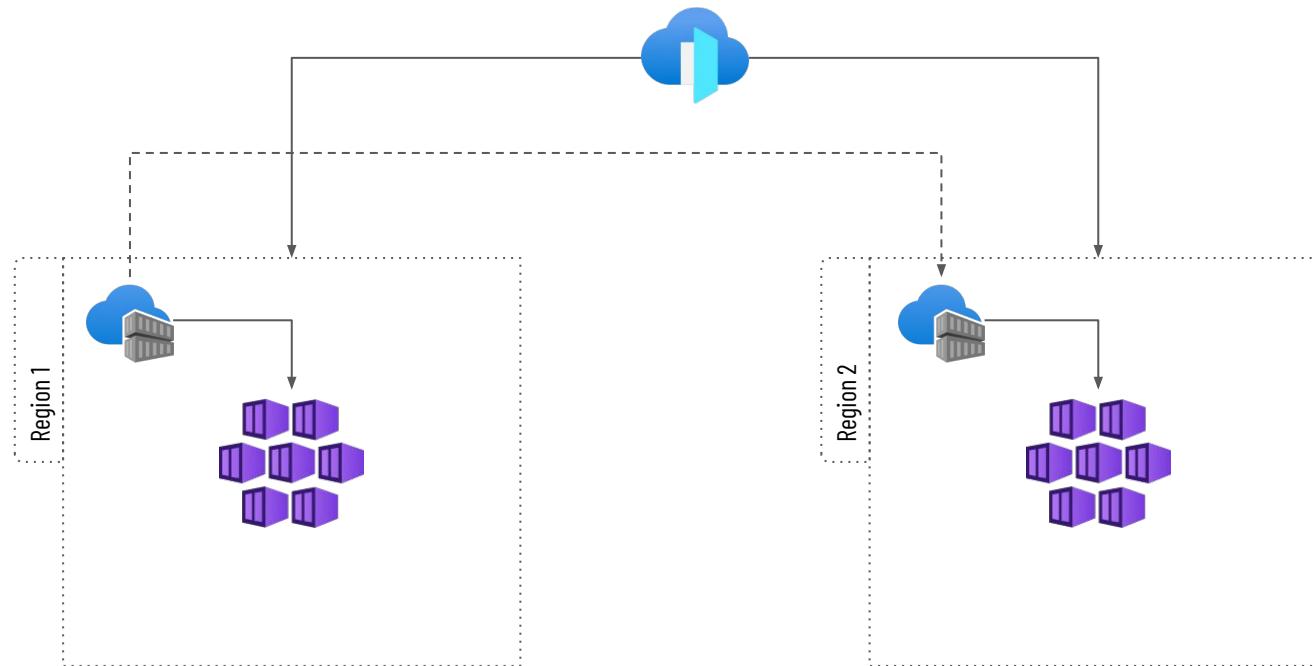
# AKS and cross-region HA



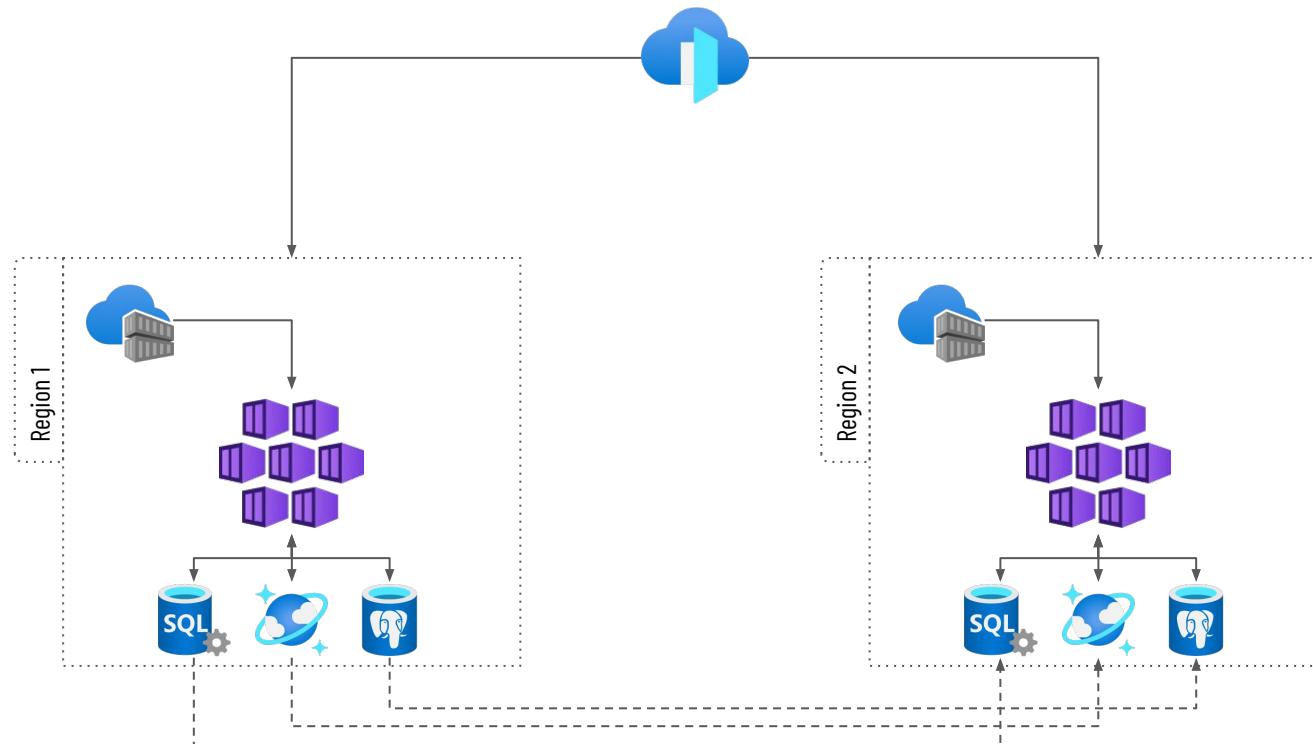
# AKS and cross-region HA



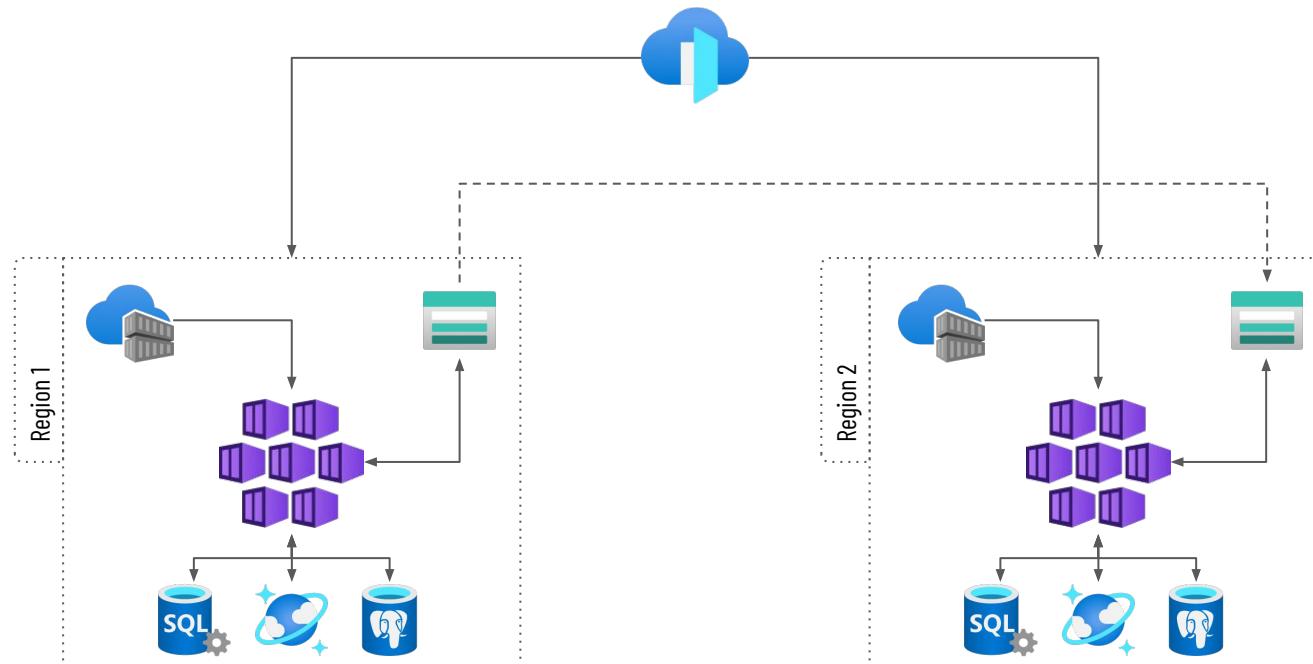
# AKS and cross-region HA



# AKS and cross-region HA



# AKS and cross-region HA

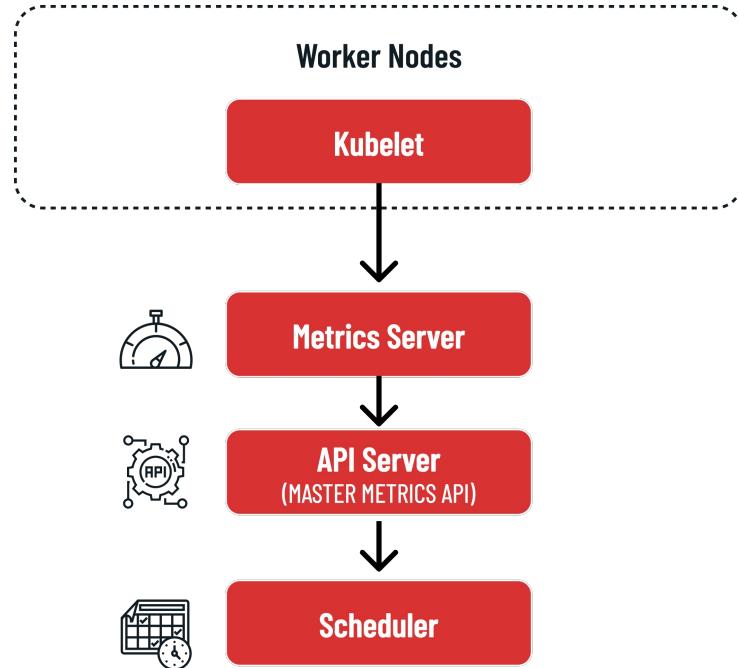


# Cluster autoscaler

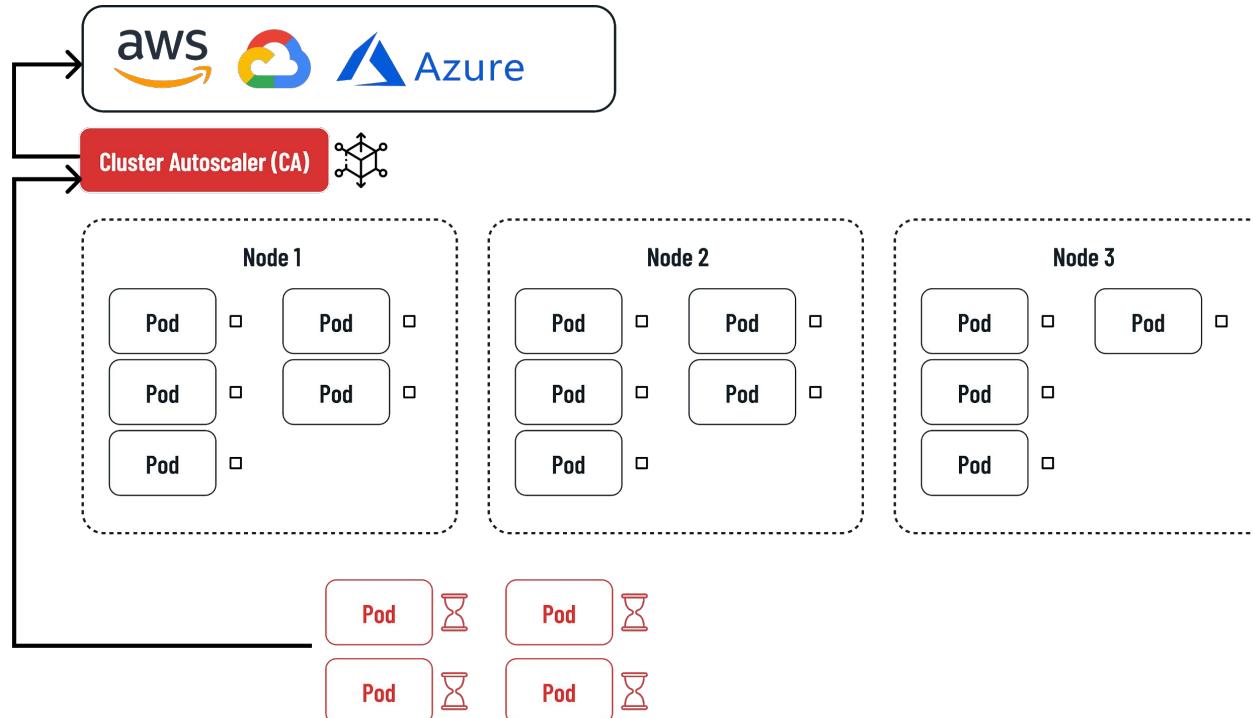
```
az aks create \
--resource-group demo-rg \
--name demo-cluster \
--enable-cluster-autoscaler \
--min-count 1 \
--max-count 3
```



# Scheduling pod



# Cluster autoscaler





That's all folks!

# Grazie!

- Il materiale sarà online nei prossimi giorni su  
<http://www.communitydays.it>



## ALESSANDRO MELCHIORI

Founder & Software developer @ CodicePlastico  
[alessandro@codiceplastico.com](mailto:alessandro@codiceplastico.com)  
[@amelchiori](https://twitter.com/amelchiori)





Grazie!

FOLLOW US

