

---

---

# Kubernetes Networking 101

AbdulKarim

---

---

# Introduction

DevOps/SRE/Platforms

[abdulkarim@relambda.com](mailto:abdulkarim@relambda.com)

@abdulkarim\_me

## Agenda

- Kubernetes - brief intro
- Pods
  - Communication Using Pod IP
  - Multi Node
  - Overlay Networks
- Services
  - Communication Using Service IP
  - IPTables
- QnA

K8s is

Open Source

Container Orchestrator

AutoScaling

Self Healing

## PODS

- There is no such thing as a **container**
- **Pod** is a group of one/many **containers**
  - Guarded by **CGroups**
  - Isolated by **Namespaces**
- Pods are **Atomic**
- Pods are **Ephemeral**

**Every Pod**  
**Gets Its**  
**Own Unique IP**  
**(IP-Per-Pod)**

**All Containers  
In A Pod Share The  
Same Network Namespace**

Vm 1  
192.168.2.3

Pod A  
192.168.0.10



Python  
port: 3000

Pod C  
192.168.0.11



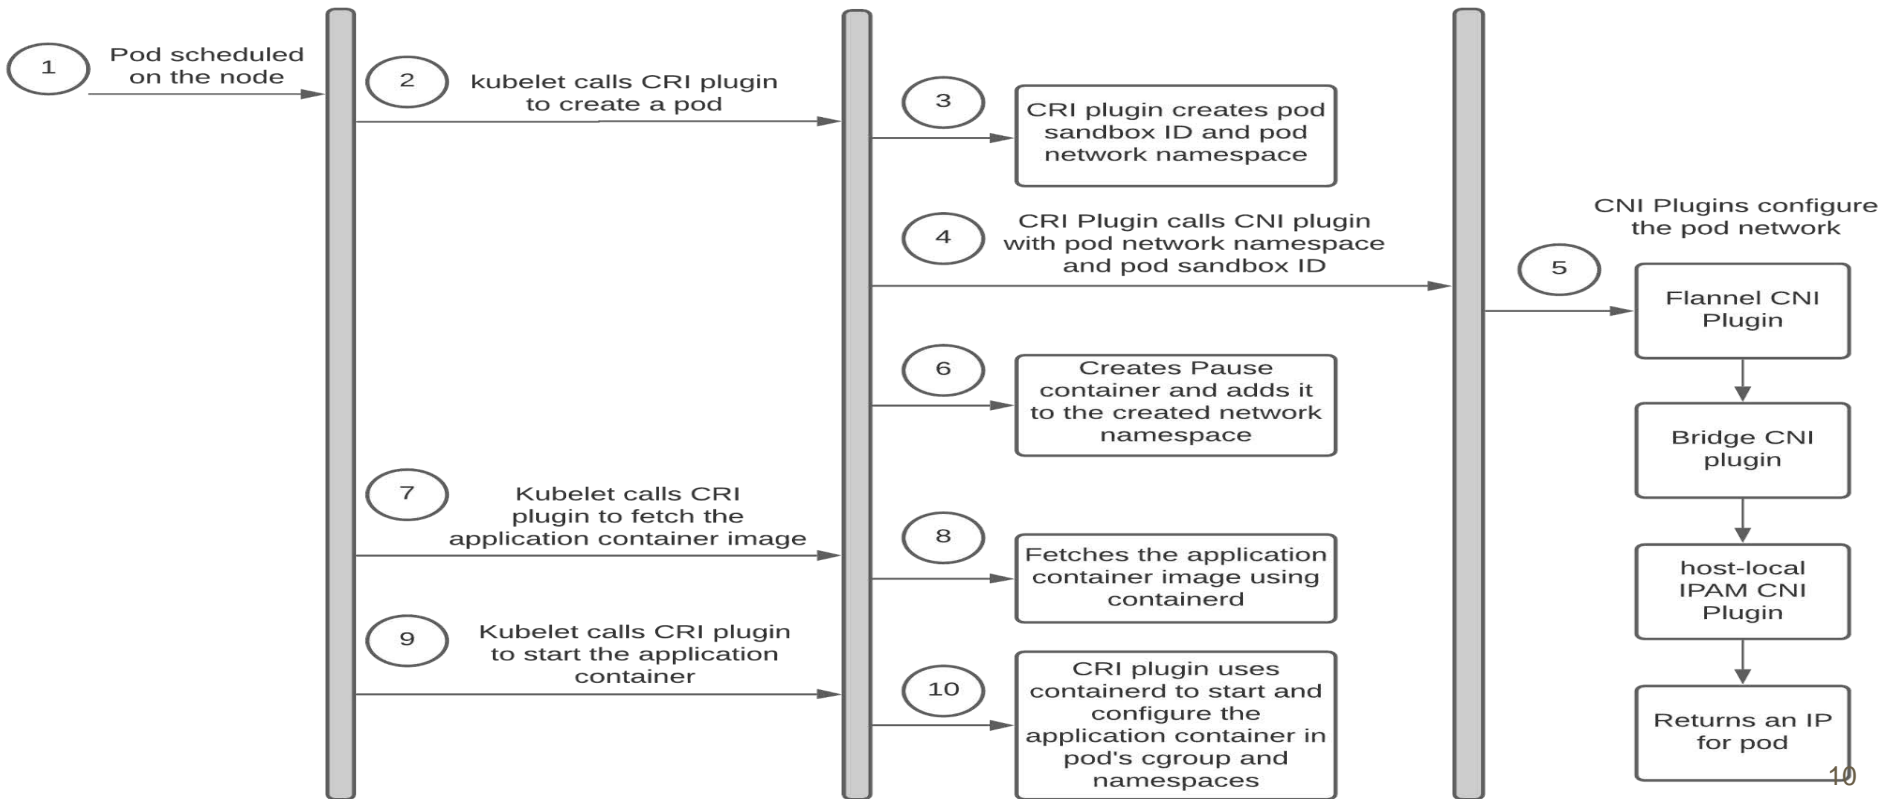
Java  
port: 3000




DataDog  
port: 8126



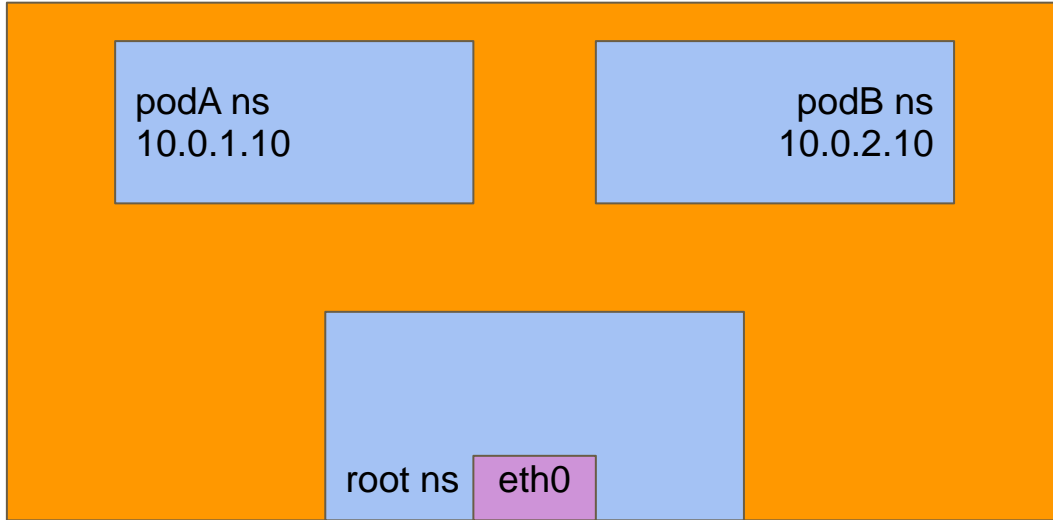
# How Does A Pod Get An IP Address?



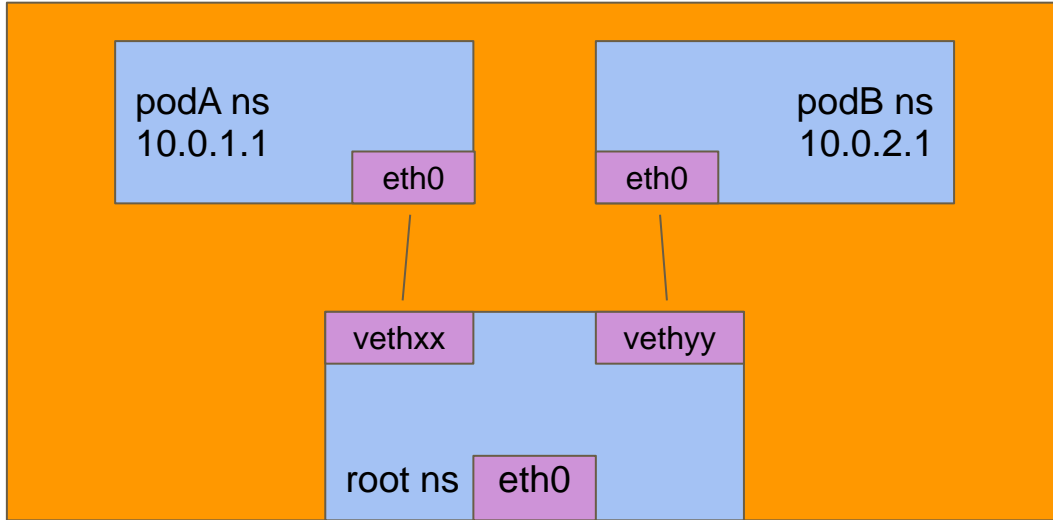
# How Does A Pod Talk To Another Pod using Pod-IP?



Node1 (10.0.240.10)

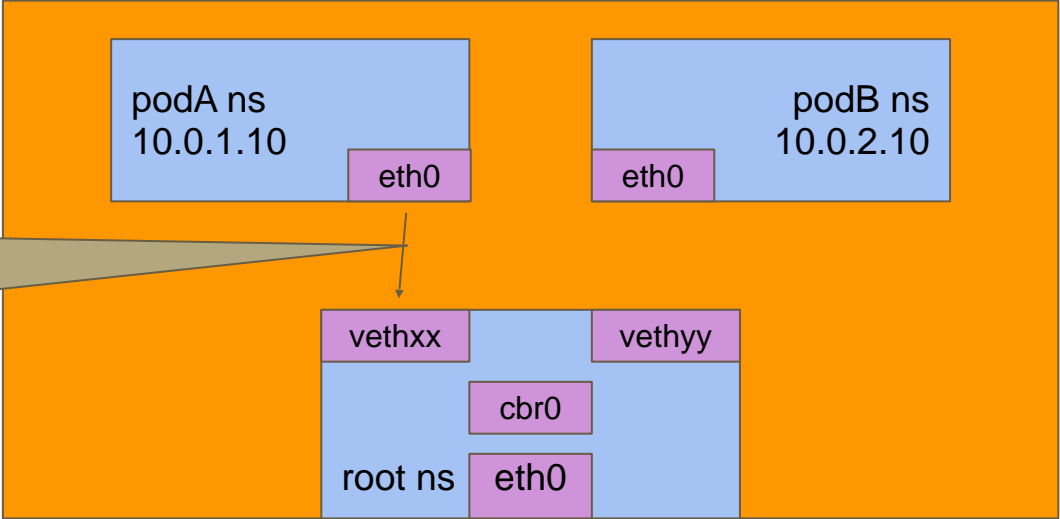


Node1 (10.0.240.10)

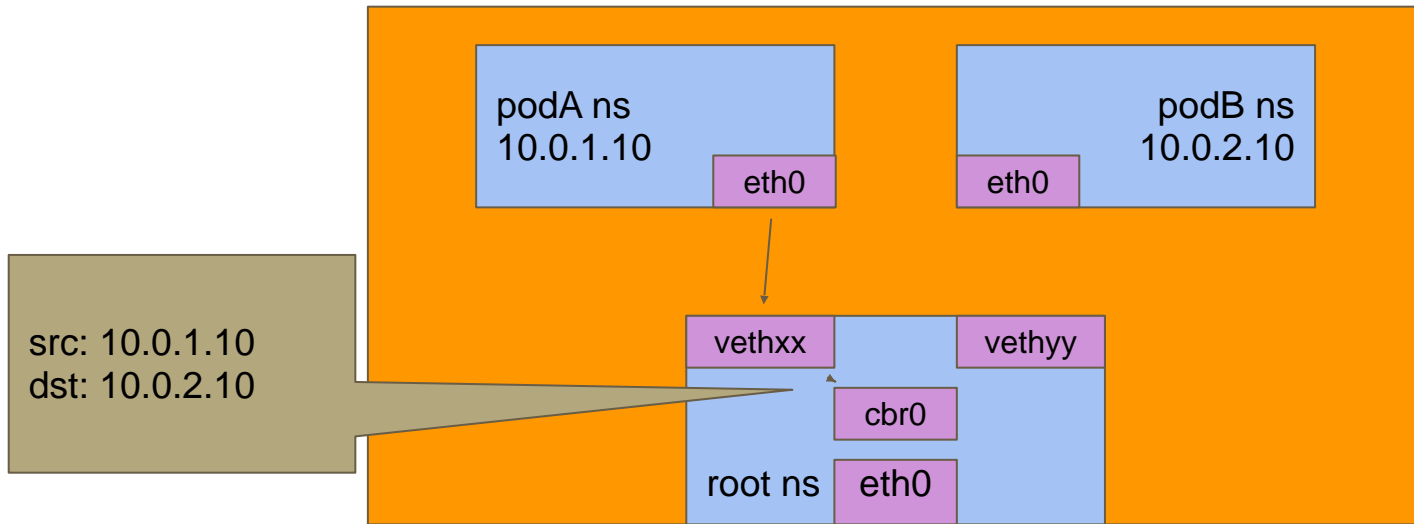


Node1 (10.0.240.1)

src: 10.0.1.10  
dst: 10.0.2.10

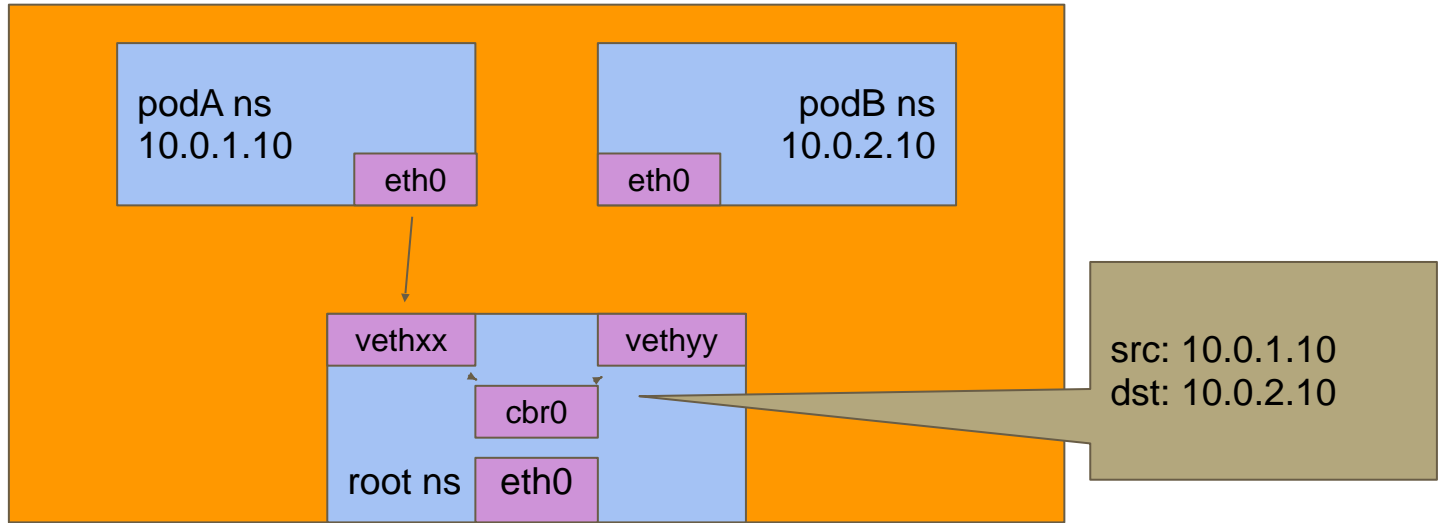


Node1 (10.0.240.10)

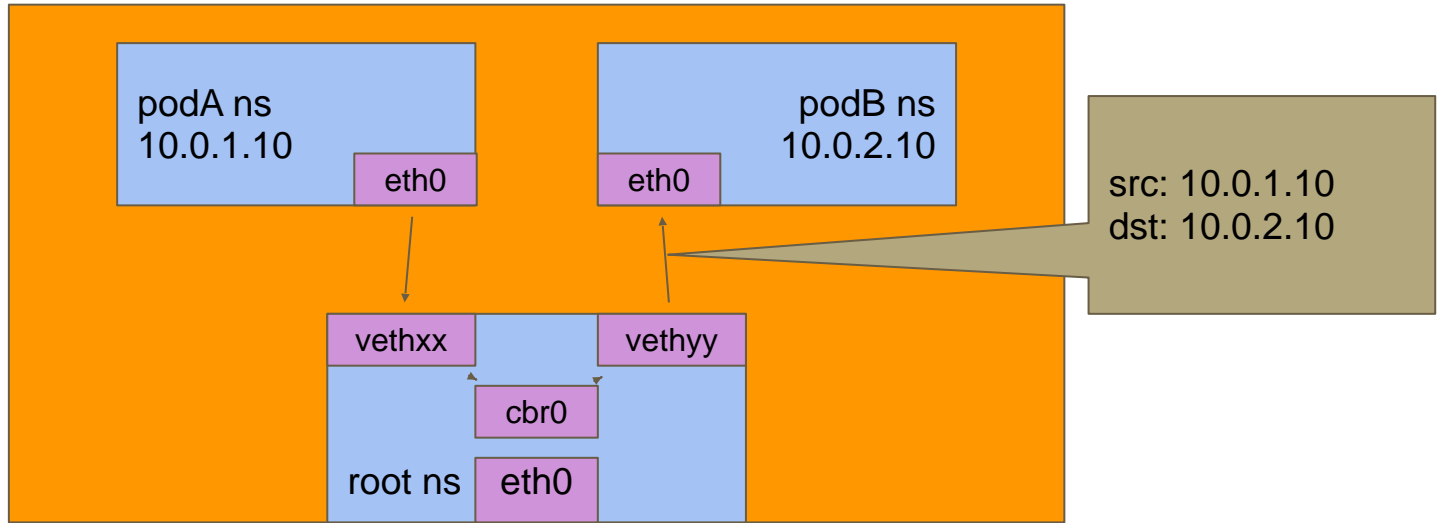


Node1 (10.0.240.10)



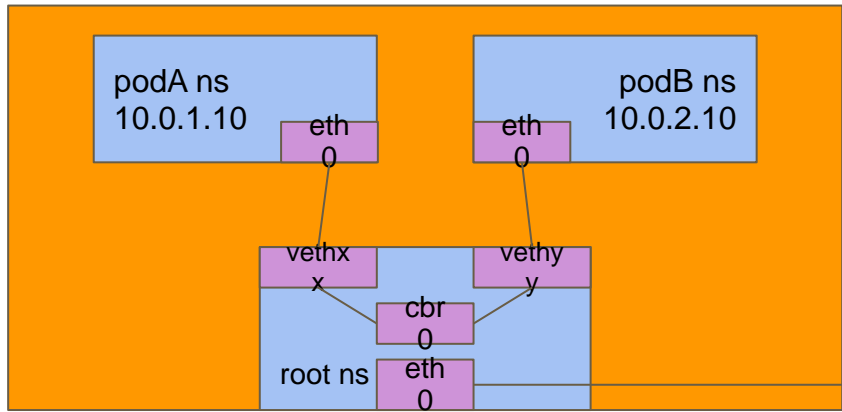


Node1 (10.0.240.10)




Node1 (10.0.240.10)

# How Does A Pod Talk To Another Pod using Pod-IP?

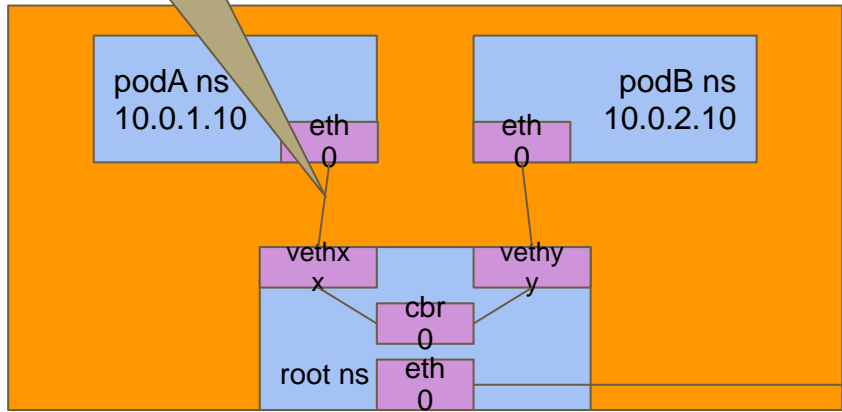


Node1 (10.0.240.10)




Node2 (10.0.250.10)

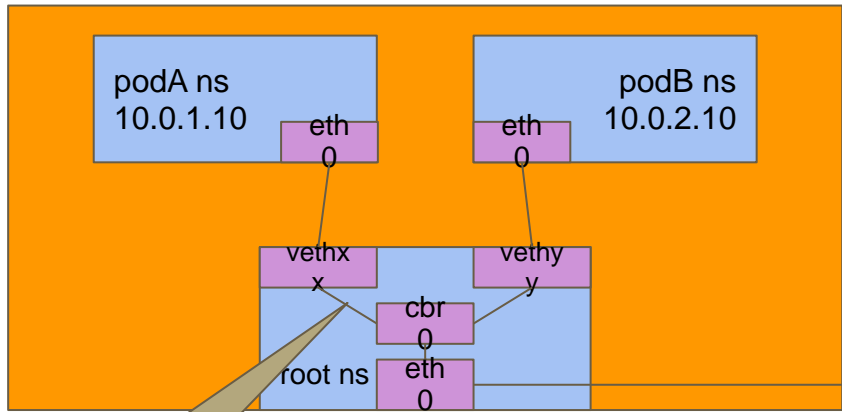
src: podA  
dst: podC



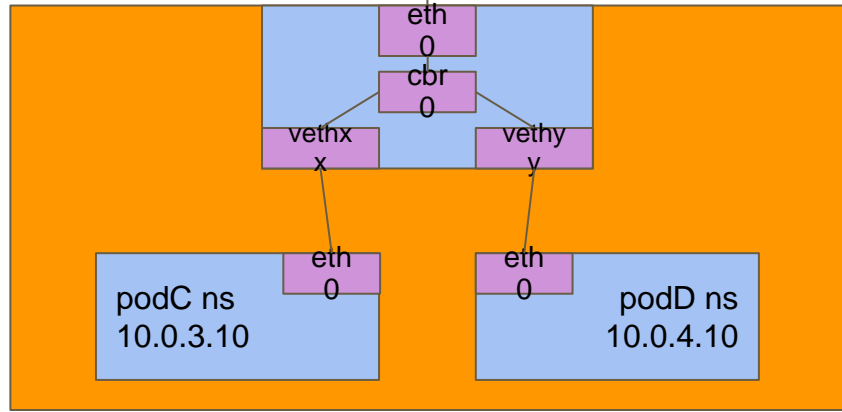
Node1 (10.0.240.10)

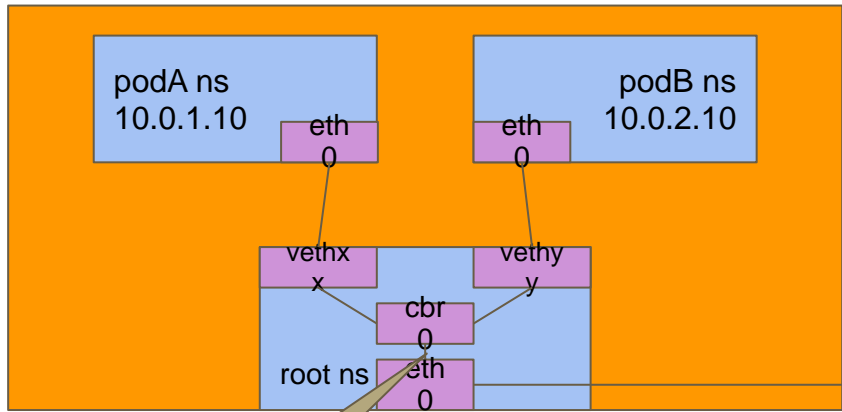


Node2 (10.0.250.10)




src: podA  
dst: podC



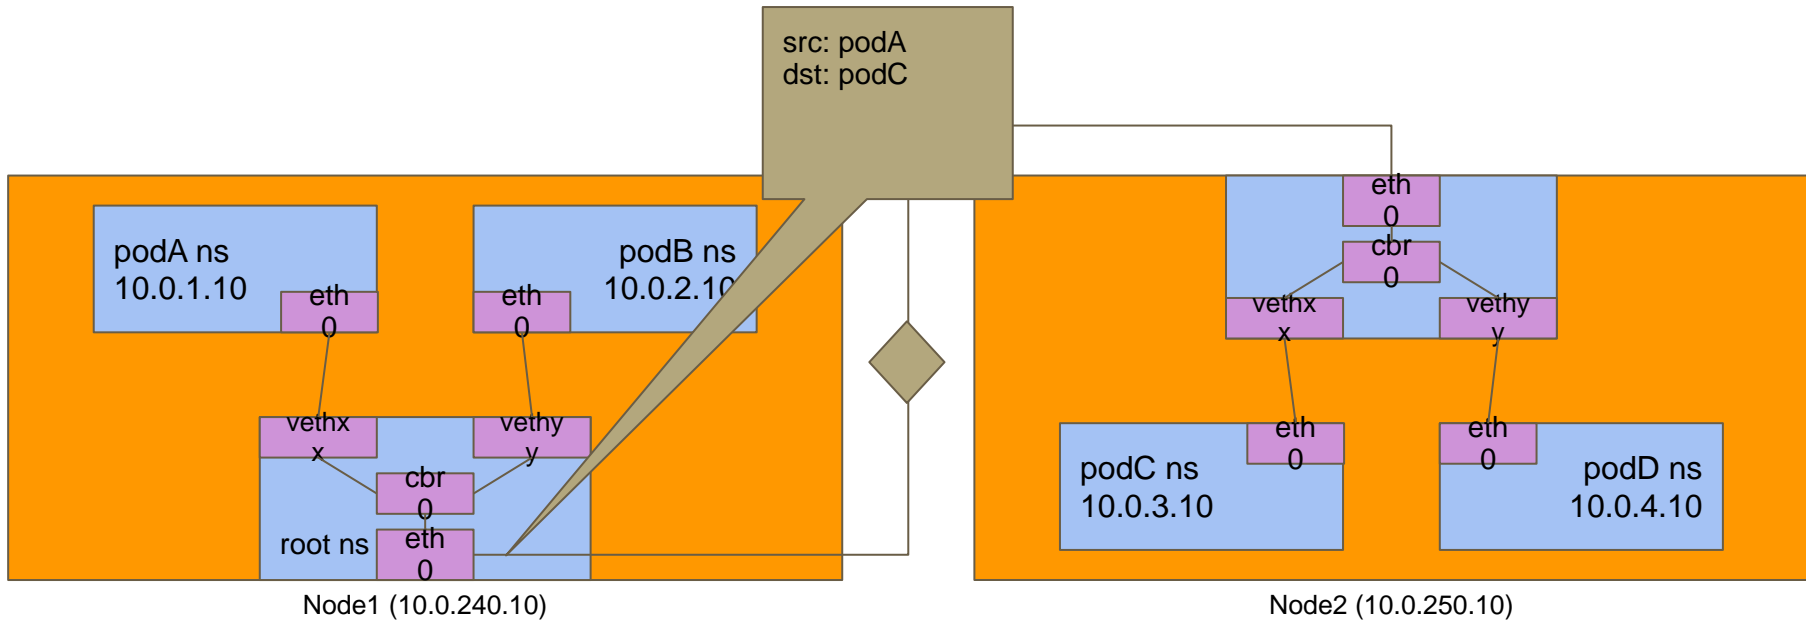


Node1 (10.0.240.10)


src: podA  
dst: podC



Node2 (10.0.250.10)







| <input checked="" type="checkbox"/> | Name         | Instance ID | Instance state | Instance type | Private IP address | Security group name                          | Launch time               | VPC ID         |
|-------------------------------------|--------------|-------------|----------------|---------------|--------------------|--|---------------------------|----------------|
| <input checked="" type="checkbox"/> | node-group-1 | i-0bce0970  | Running        | m5.2xlarge    | 192.150.2.21       | eks-2LgdDngQ-node-20230611092714934500000008 | 2023/07/13 19:58 GMT+5:30 | vpc-0fae52340a |

### Instance: i-041c01e00bce0970 (sss-prod-node-group-1)

Public IPv4 address

-

Public IPv4 DNS

-

Subnet ID

subnet-03538cbee0b82718a (eks-vpc-private-ap-south-1b)

Private IPv4 addresses

- 192.150.2.21
- 192.150.2.187

Private IP DNS name (IPv4 only)

ip-192-150-2-21.ap-south-1.compute.internal

IPv6 addresses

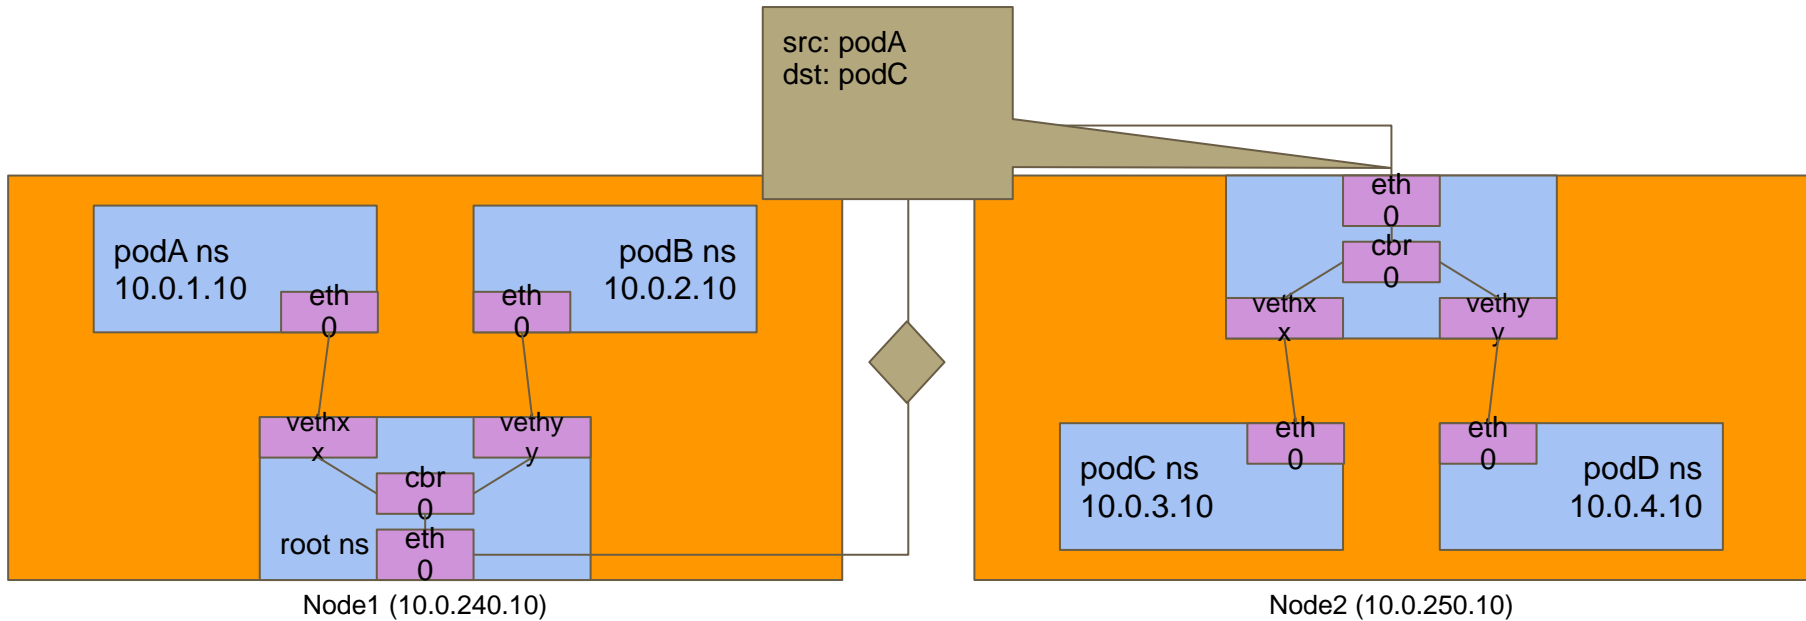
-

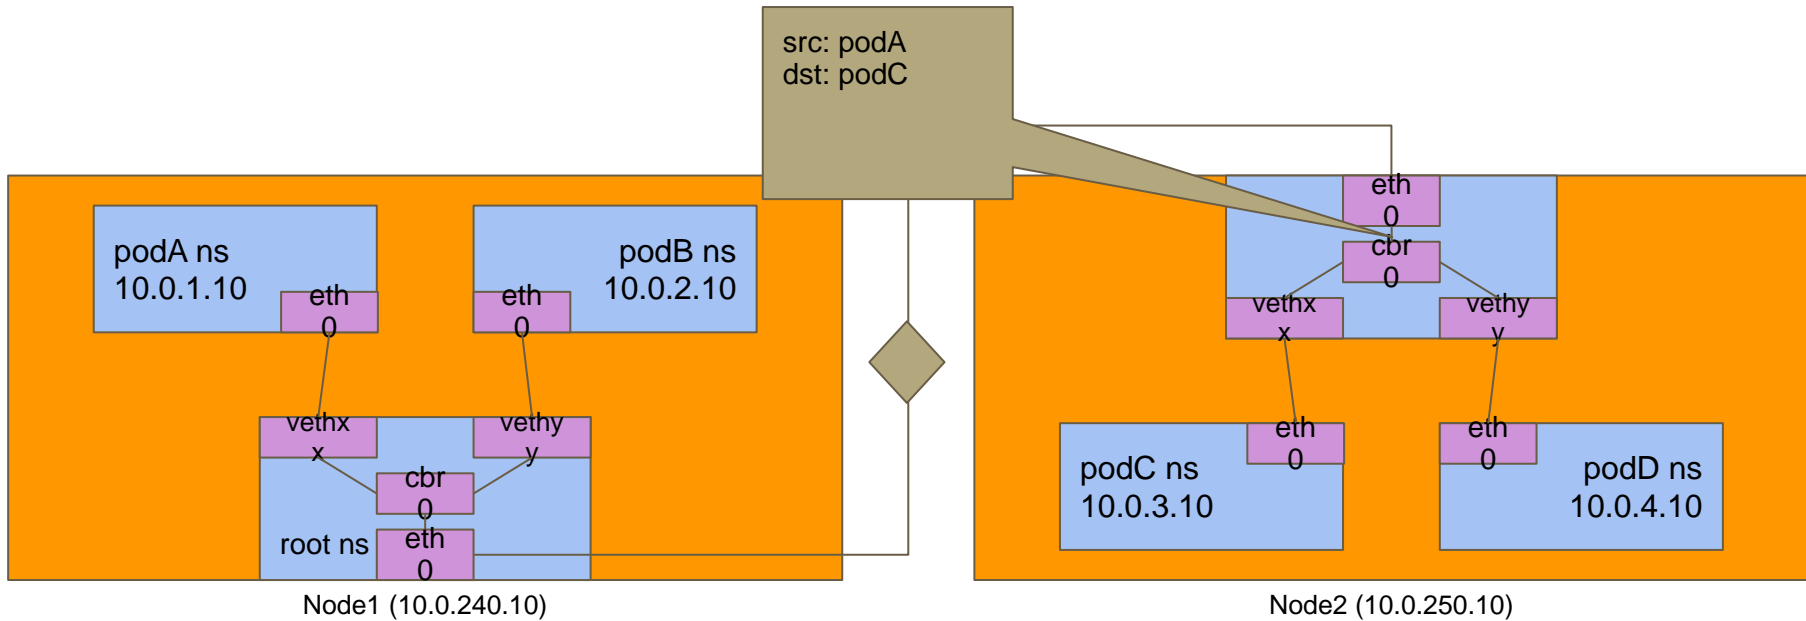
VPC ID

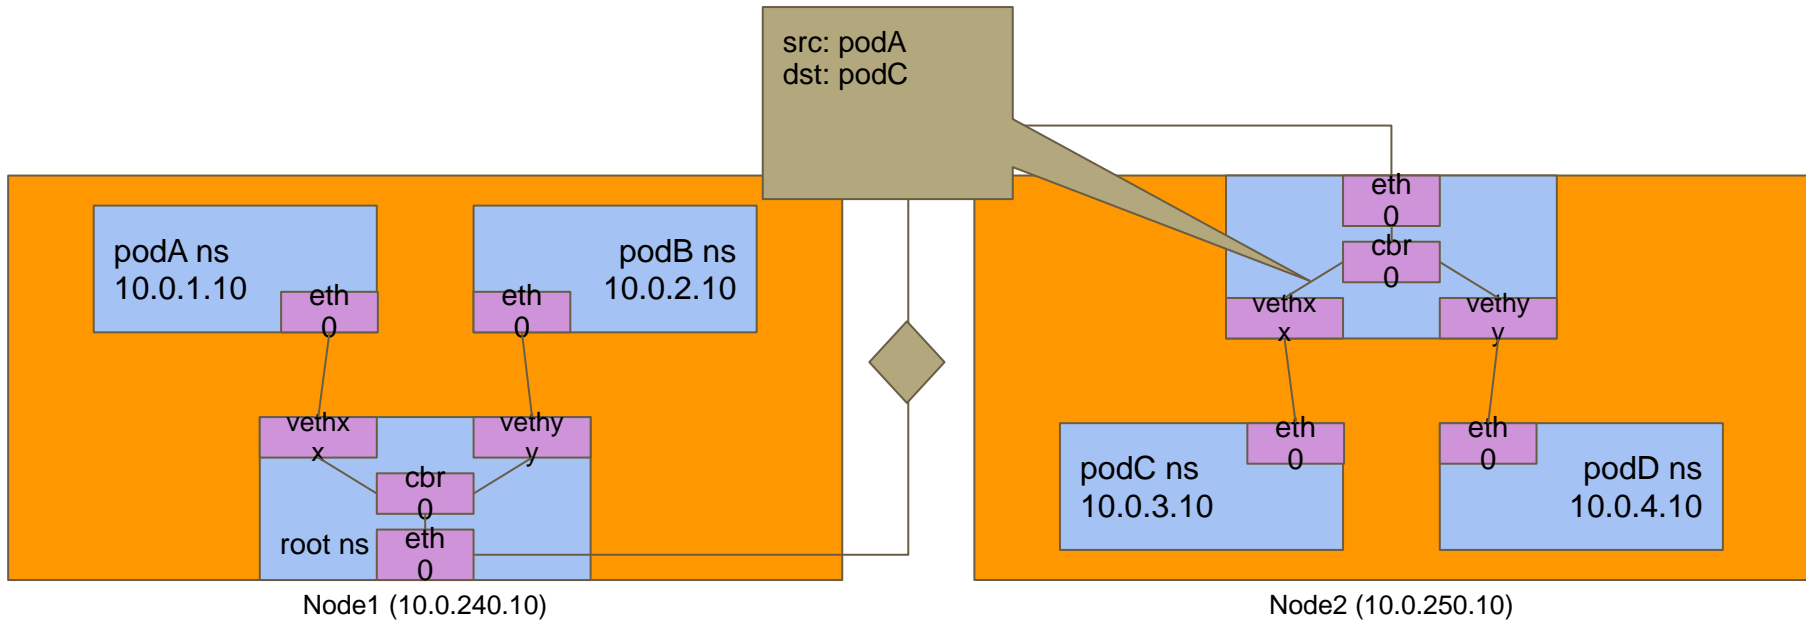
vpc-0fae52340ba89ac8a (eks-vpc)

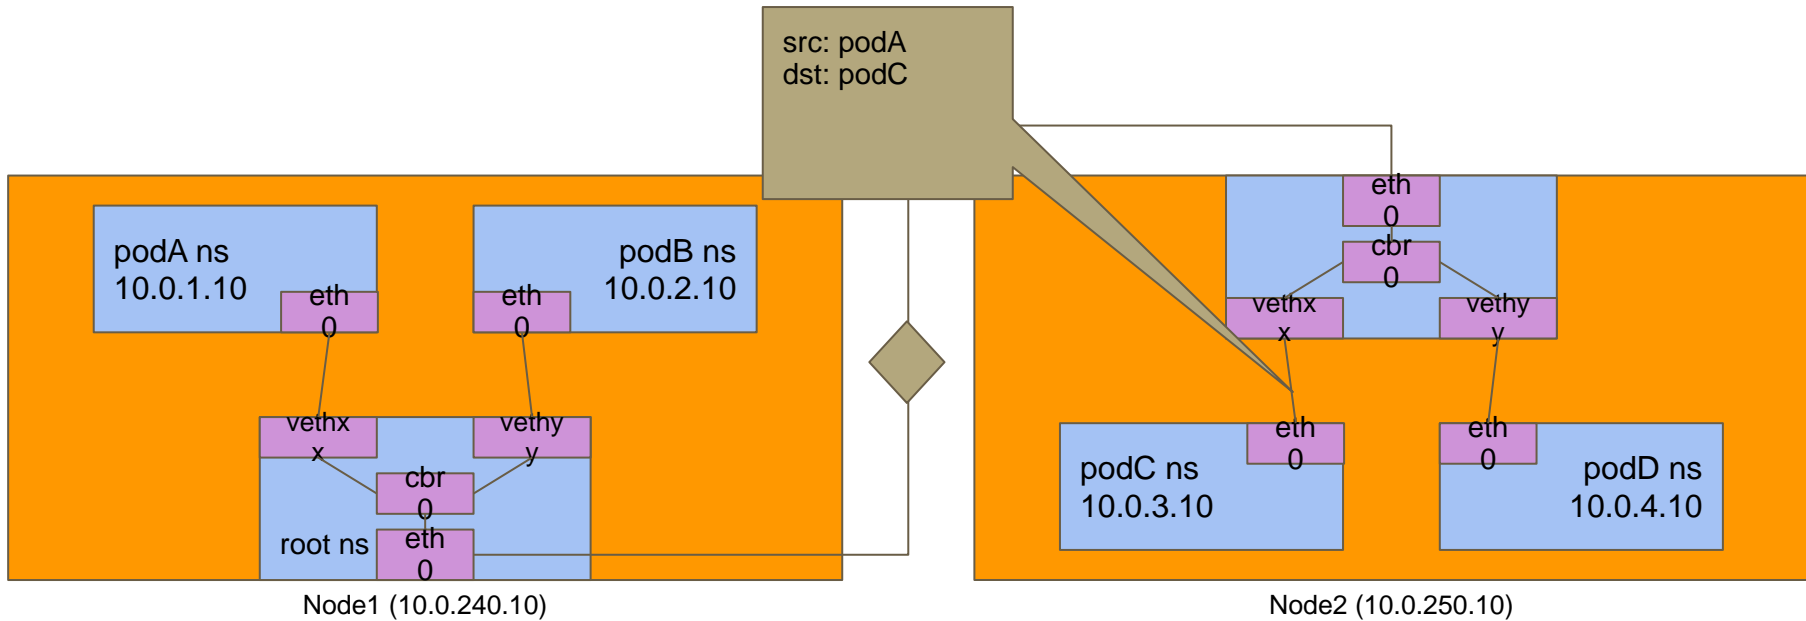
Secondary private IPv4 addresses

- 192.150.2.182
- 192.150.2.50
- 192.150.2.83
- 192.150.2.190
- 192.150.2.123
- 192.150.2.217
- 192.150.2.166
- 192.150.2.164
- 192.150.2.67
- 192.150.2.174
- 192.150.2.239
- 192.150.2.76
- 192.150.2.235
- 192.150.2.72
- 192.150.2.81
- 192.150.2.92
- 192.150.2.93
- 192.150.2.88
- 192.150.2.120



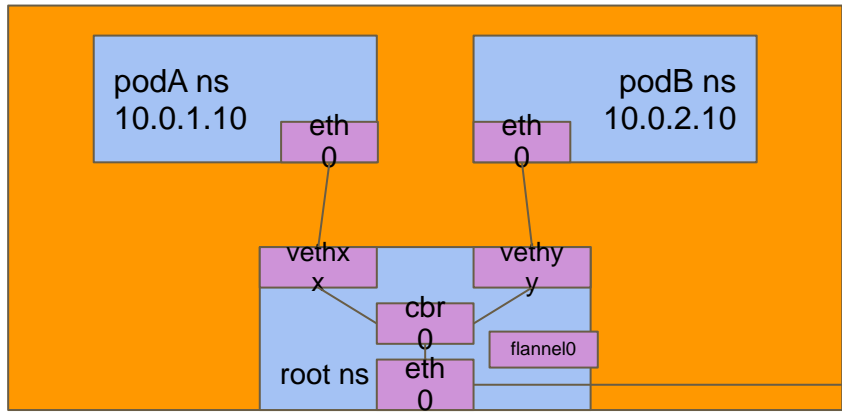




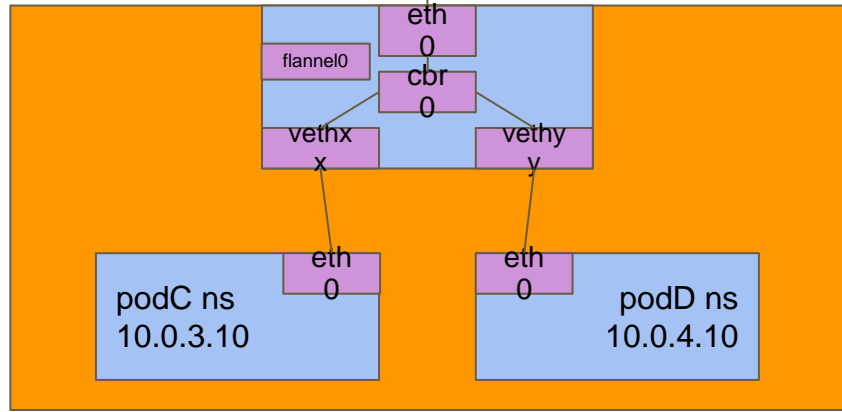


# Overlay Network

**IP Ranges Gets Exhausted**



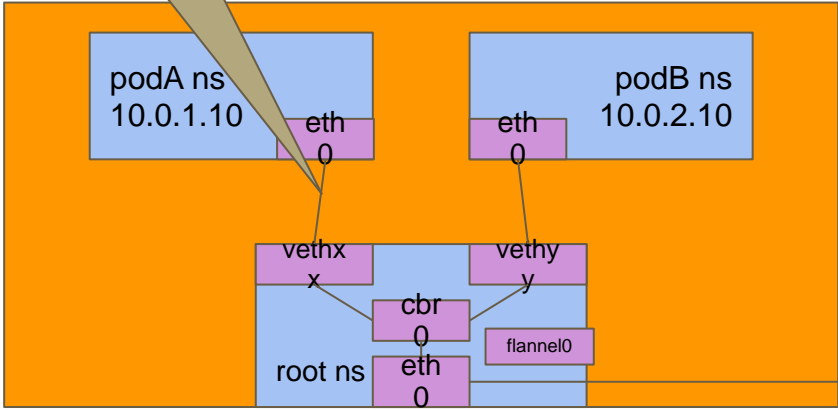
Node1 (10.0.240.10)



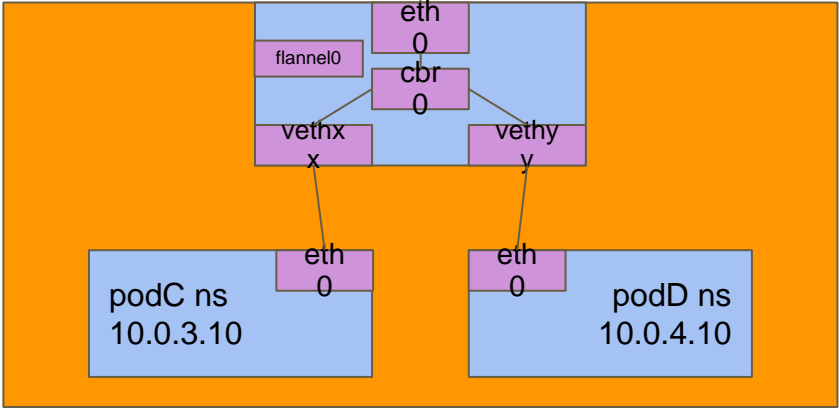
Node2 (10.0.250.10)



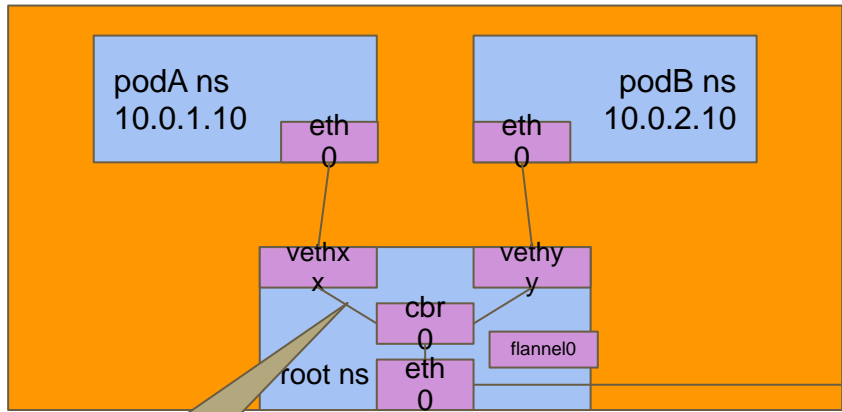
src: podA  
dst: podC



Node1 (10.0.240.10)




Node2 (10.0.250.10)



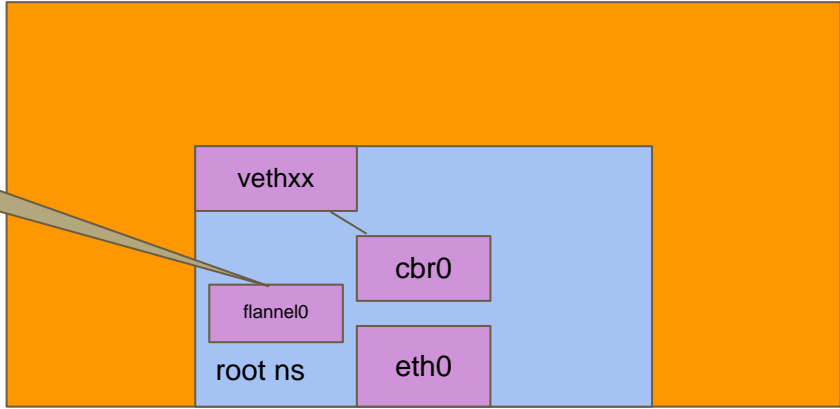
Node1 (10.0.240.10)

src: podA  
dst: podC

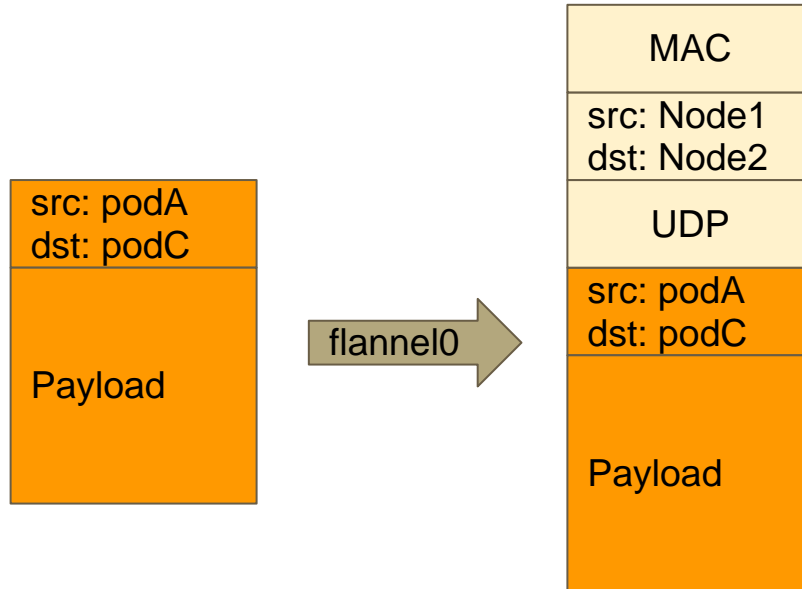


Node2 (10.0.250.10)

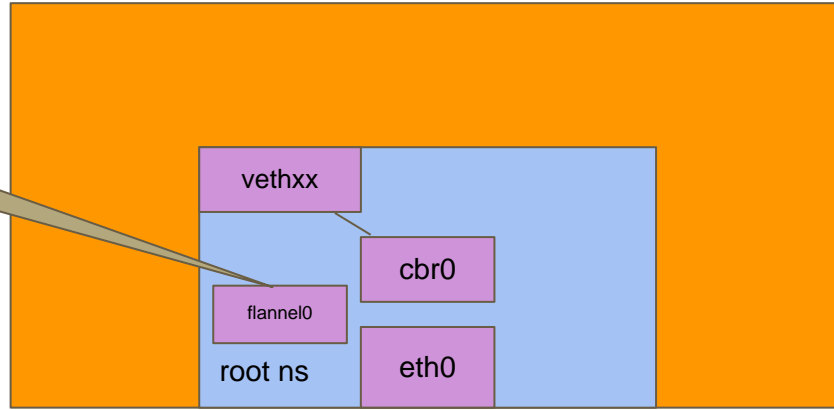
src: podA  
dst: podC



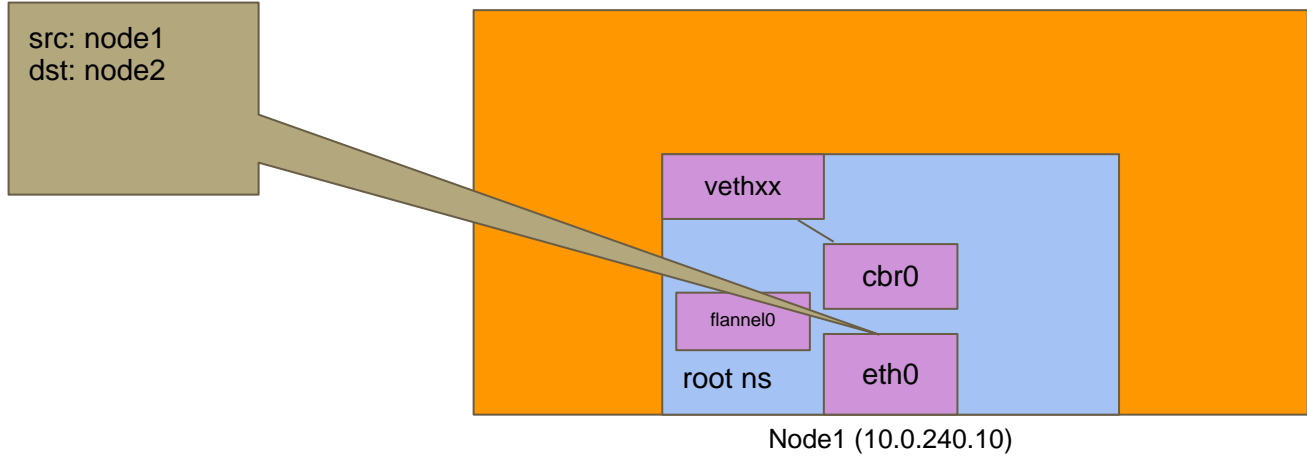
Node1 (10.0.240.10)

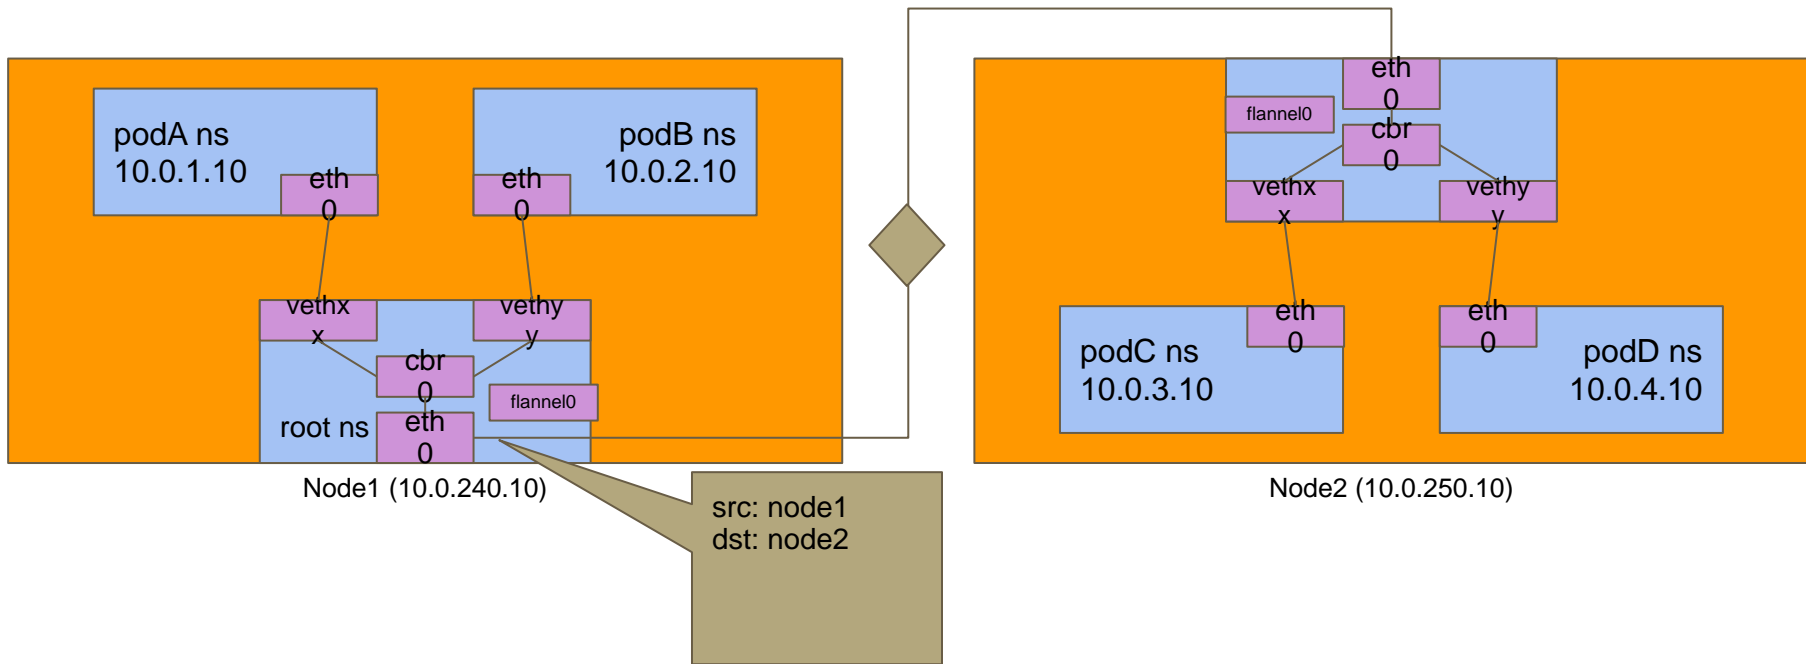



src: podA  
dst: podG  
src: node1  
dst: node2




Node1 (10.0.240.10)






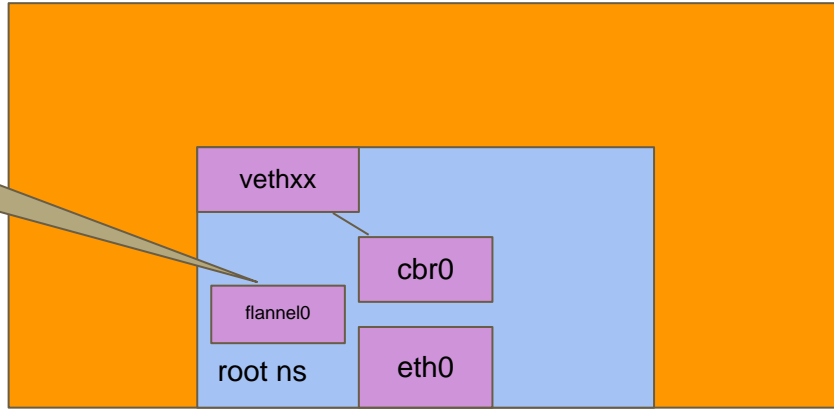







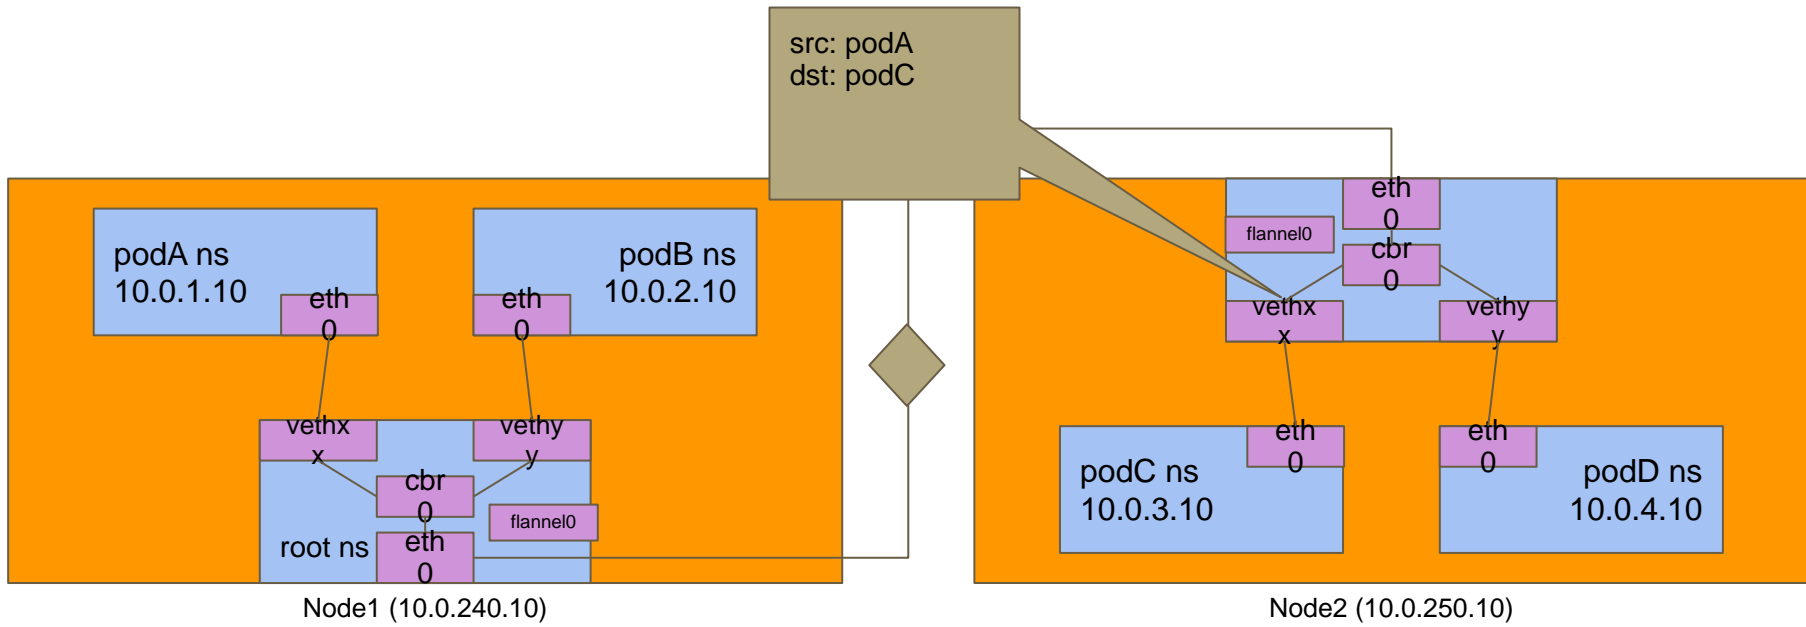


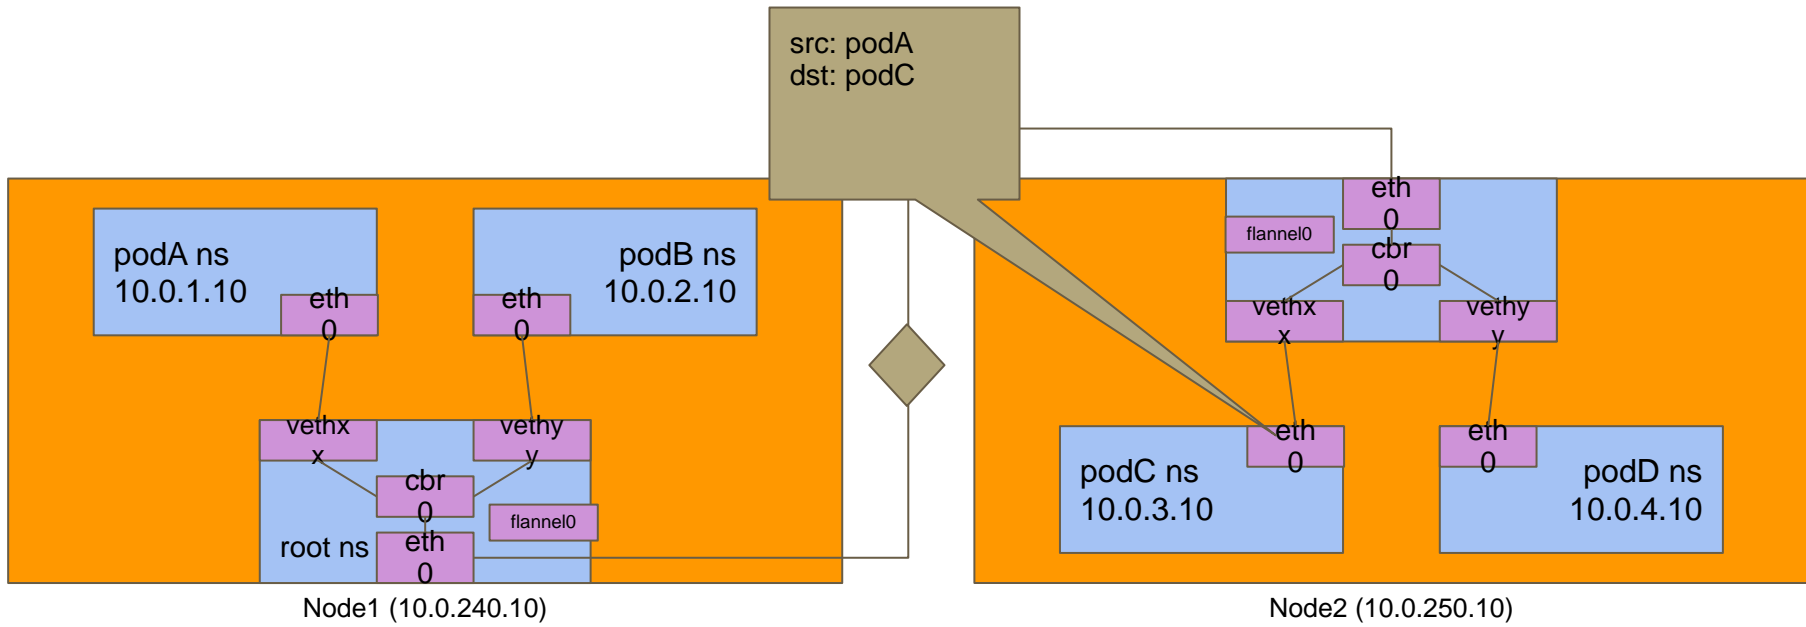
src: node1  
dst: node2  
src: podA  
dst: podC



Node2 (10.0.250.10)







What **If**  
The Pod **IP**  
Address **Changes?**

**Service**

Abstraction

For A

**Group Of Pods**



apiVersion: v1

kind: Service

metadata:

name: hello-kubernetes

spec:

ports:

- port: 80

protocol: TCP

targetPort: 8080

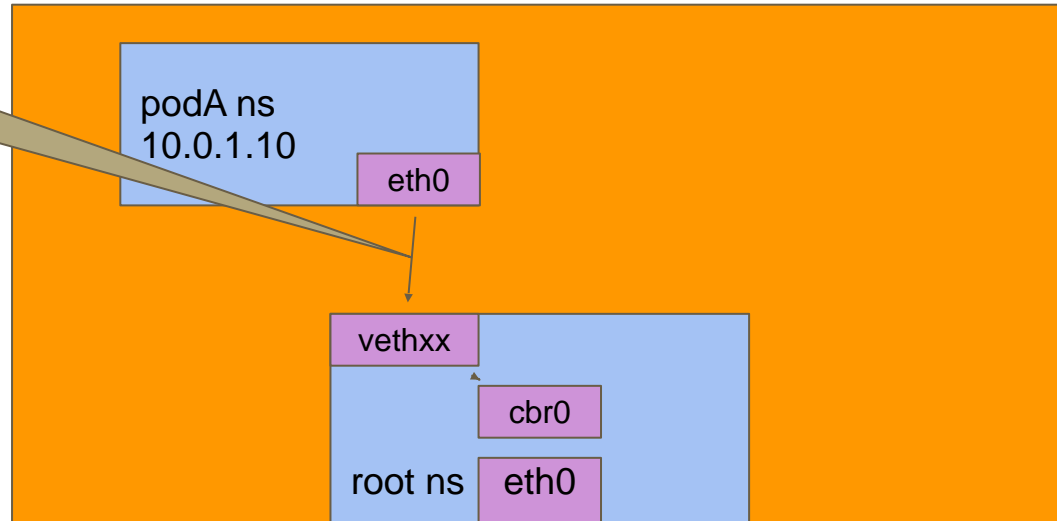
selector:

**app: hello-kubernetes**

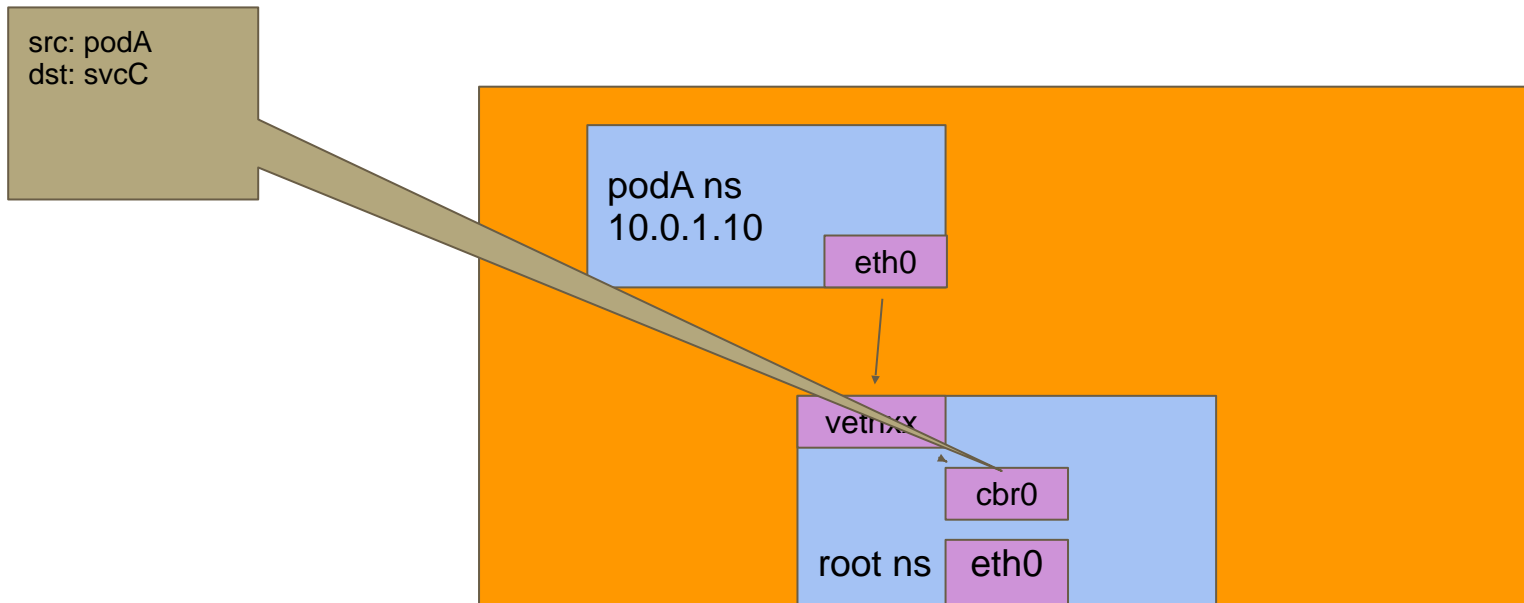
type: ClusterIP

# How Does A Pod Talk To Another Pod Using Service IP?

src: podA  
dst: svcC



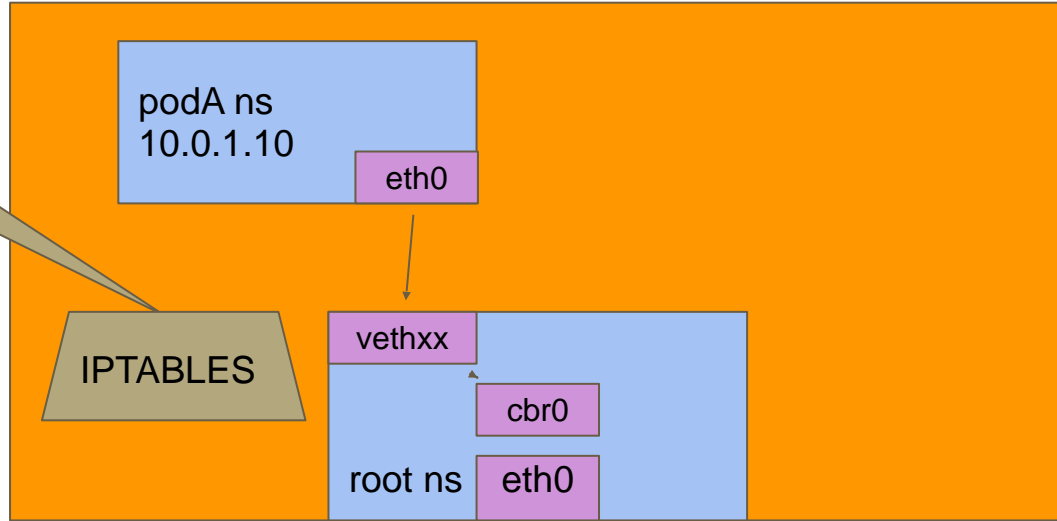
Node1 (10.0.240.10)



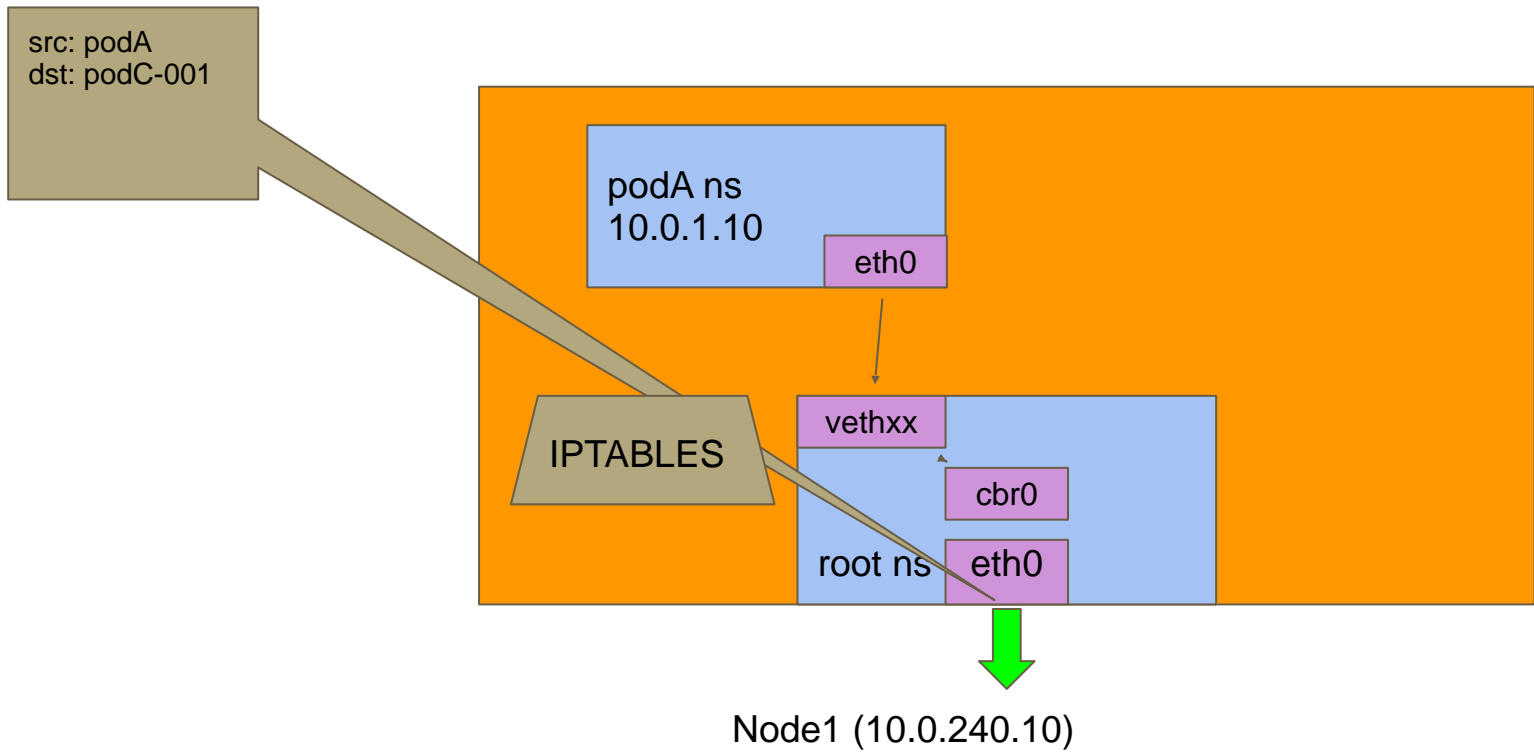
Node1 (10.0.240.10)

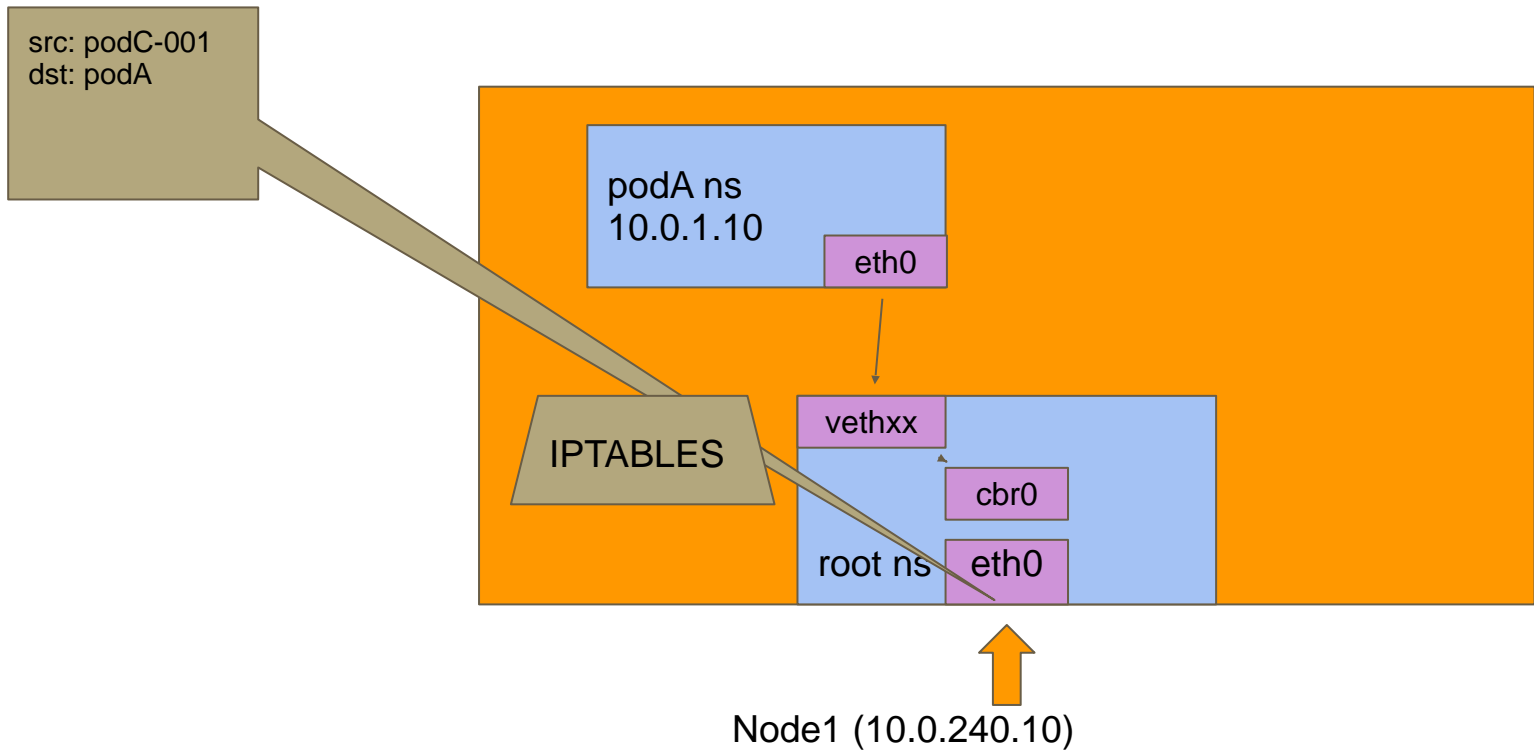
src: podA  
dst: sveG  
dst: podC-001

DNAT,  
conntrack

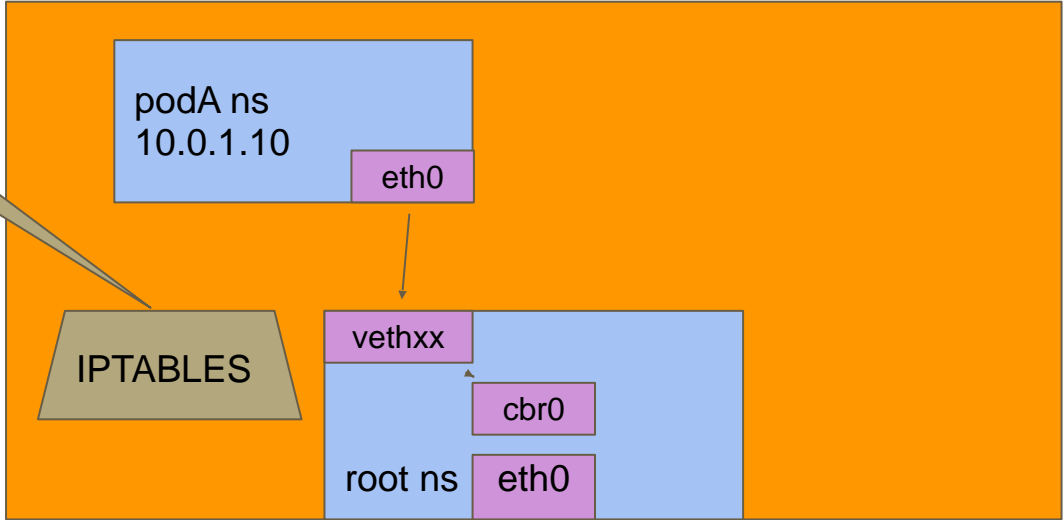


Node1 (10.0.240.10)






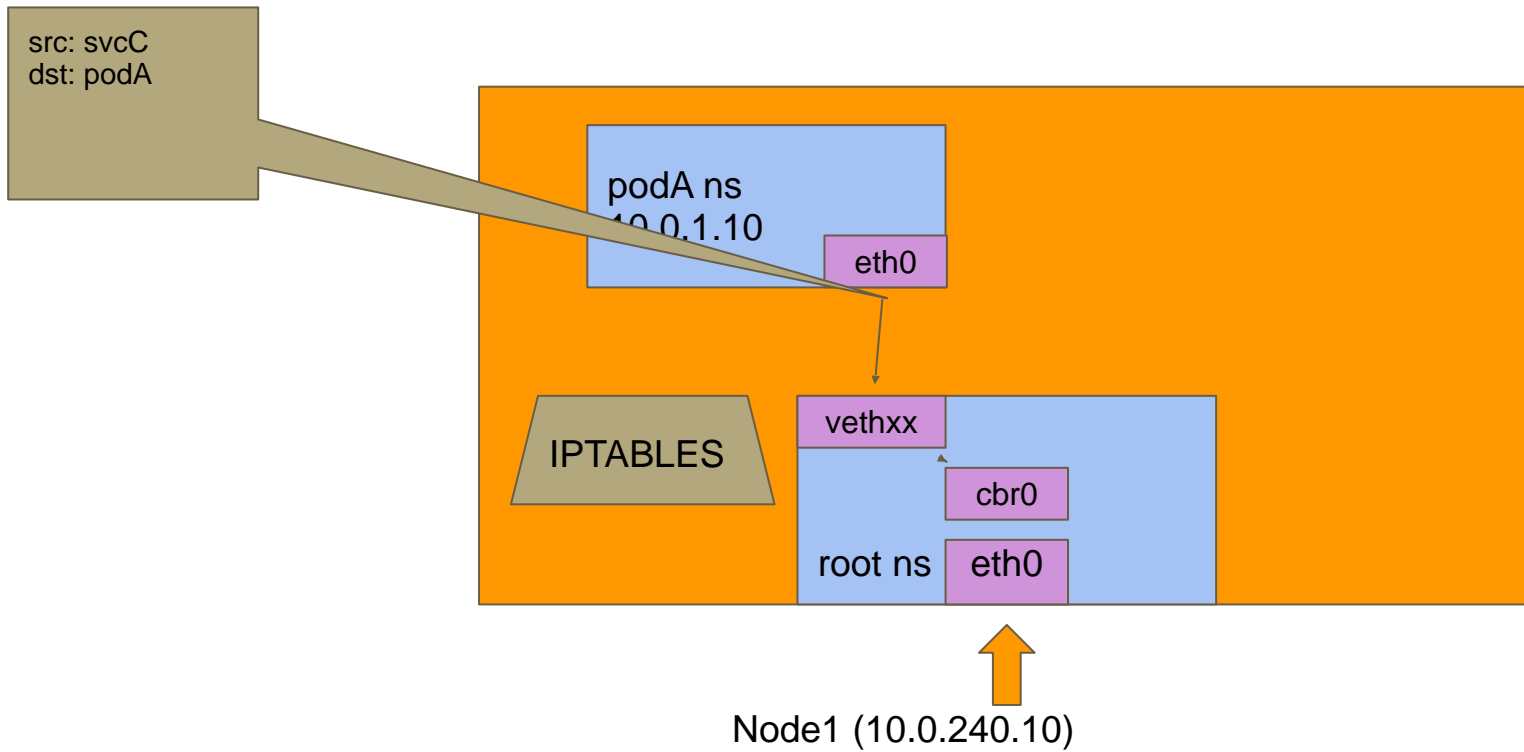
src: podC-001  
src: svcC  
dst: podA



Node1 (10.0.240.10)







QnA

THANK YOU!

# Resources

<https://learncloudnative.com/blog/2023-05-31-kubeproxy-iptables>

<https://www.redhat.com/sysadmin/kubernetes-pods-communicate-nodes>

<https://www.youtube.com/watch?v=0Omvgd7Hg1I>

<https://www.youtube.com/watch?v=lnZVNuKY5GY&t=867s>