

[Network] K8S Overlay Network (IPIP -> VXLAN)

K8S Overlay Network

IPIP -> VXLAN

) POD가
(pod 가)

Calico IP-IP Network VXLAN

Node : Controller / Worker01 / Worker02

```
## Controller
# Mode      IPIPMode
calicctl get ippool -o wide
NAME          CIDR          NAT      IPIPMode
VXLANMode    DISABLED     DISABLEBGPEXPORT   SELECTOR
default-ipv4-ippool  192.168.0.0/16  true    Always    Never
false        false
all()

#      Manifest YAML
kubectl delete -f calico.yml
```

```
## Controller / Worker
# 가          tunl0      가
sudo rm -rf /var/run/calico/
sudo rm -rf /var/lib/calico/
sudo rm -rf /etc/cni/net.d/
sudo rm -rf /var/lib/cni/
sudo reboot
```

```
## Controller
# Manifest. calico.yaml      VXLAN
```

```

livenessProbe:
  exec:
    command:
      - /bin/calico-node
      - -felix-live
      # - -bird-live          // VXLAN    bird(BGP)

    periodSeconds: 10
    initialDelaySeconds: 10
    failureThreshold: 6
    timeoutSeconds: 10
readinessProbe:
  exec:
    command:
      - /bin/calico-node
      - -felix-ready
      # - -bird-ready     //

# Enable IPIP
- name: CALICO_IPV4POOL_IPIP
  value: "Never"           // Always --> Never

# Enable or Disable VXLAN on the default IP pool.
- name: CALICO_IPV4POOL_VXLAN
  value: "Always"          // Never --> Always

kind: ConfigMap
apiVersion: v1
metadata:
  name: calico-config
  namespace: kube-system
data:
  # Typha is disabled.
  typha_service_name: "none"
  # Configure the backend to use.
  calico_backend: "vxlan"        // "bird" --> "vxlan"
  .
  #
# kubectl apply -f calico.yaml

```

```

# Calico Node . Ready
kubectl get nodes -o wide -A

# Calico Pod . kube-system PoD      가
kubectl get pod -o wide -A

# Calico Type . BIRD
sudo calicectl node status
Calico process is running.
The BGP backend process (BIRD) is not running.

# Network     VXLANMODE 가
calicectl get ippool -o wide
NAME          CIDR          NAT      IPIP MODE
VXLANMODE    DISABLED    DISABLEBGPEXPORT  SELECTOR
default-ipv4-ippool  192.168.0.0/16  true    Never
Always        false       false

#           tunl0      가      vxlan      가
#           vxlan

hostway@controller:~$ route -n
Kernel IP routing table
Destination     Gateway         Genmask        Flags Metric
Ref  Use Iface
0.0.0.0         10.10.10.1   0.0.0.0        UG    0      0
0 ens18
10.10.10.0     0.0.0.0       255.255.255.0  U      0      0
0 ens18 // External (SNAT)
172.17.0.0      0.0.0.0       255.255.0.0    U      0      0
0 docker0 // Container Runtime Bridge
192.168.5.0     192.168.5.0   255.255.255.192 UG    0      0
0 vxlan.calico // Worker01
192.168.30.64   192.168.30.64  255.255.255.192 UG    0      0
0 vxlan.calico // Worker02
192.168.49.0    0.0.0.0       255.255.255.192 U      0      0
0 *             // Controller  vxlan
192.168.49.1    0.0.0.0       255.255.255.255 UH    0      0
0 cali09ae4a7064b // Node(Worker01)가      GW
192.168.49.2    0.0.0.0       255.255.255.255 UH    0      0
0 calilfdac863dc5 // Node(Worker02)가      GW

```

```
# Worker
hostway@controller:~$ ip nei | grep vxlan
192.168.5.0  dev  vxlan.calico  lladdr  66:8c:33:86:44:ce
PERMANENT
192.168.30.64  dev  vxlan.calico  lladdr  66:fb:72:20:22:a1
PERMANENT
```

```
# VXLAN Traffic Port  UDP
udp      0      0 0.0.0.0:4789          0.0.0.0:*
```

```
# PoD
```

```
hostway@controller:~$ kubectl create deployment sampleos --image=gcr.io/google-samples/kubernetes-bootcamp:v1 --replicas=3
```

```
deployment.apps/sampleos created
```

```
hostway@controller:~$ kubectl get pod -o wide
```

NAME	READY	STATUS	RESTARTS	AGE
IP	NODE	NOMINATED NODE	READINESS GATES	
sampleos-646dc9654b-8xjw9	1/1	Running	0	45s
192.168.5.11	worker01	<none>	<none>	
sampleos-646dc9654b-gxn75	1/1	Running	0	45s
192.168.5.10	worker01	<none>	<none>	
sampleos-646dc9654b-snkgxg	1/1	Running	0	45s
192.168.30.75	worker02	<none>	<none>	

```
# VXLAN
```

```
// Controller
```

```
1) worker01  worker02          POD  Ping
```

```
hostway@controller:~$ kubectl exec -it sampleos-646dc9654b-8xjw9 -- ping 192.168.30.75
PING 192.168.30.75: 56 data bytes
64 bytes from 192.168.30.75: icmp_seq=0 ttl=115 time=92.124 ms
64 bytes from 192.168.30.75: icmp_seq=1 ttl=115 time=79.735 ms
64 bytes from 192.168.30.75: icmp_seq=2 ttl=115 time=79.233 ms
```

```
2)                                     tcpdump
```

```
sudo tcpdump -i ens18 -w vxlan.pcap
```

3) Wireshark

UDP

vxlan.pcap

File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

Apply a display filter ... <Ctrl-/>

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	192.168.5.11	192.168.49.2	DNS	158	Standard query 0x475e A storage.googleapis.com.default.svc.cluster.local
2	0.000001	192.168.5.11	192.168.49.2	DNS	158	Standard query 0xdf0e AAAA storage.googleapis.com.default.svc.cluster.local
3	0.002360	192.168.49.2	192.168.5.11	DNS	251	Standard query response 0xdf0e No such name AAAA storage.googleapis.com.default.svc.cluster.local
4	0.002958	192.168.49.2	192.168.5.11	DNS	251	Standard query response 0x475e No such name A storage.googleapis.com.default.svc.cluster.local
5	0.003747	192.168.5.11	192.168.49.2	DNS	150	Standard query 0xebca A storage.googleapis.com.svc.cluster.local
6	0.003939	192.168.5.11	192.168.49.2	DNS	150	Standard query 0x7f90 AAAA storage.googleapis.com.svc.cluster.local

> Frame 2: 158 bytes on wire (1264 bits), 158 bytes captured (1264 bits)
> Ethernet II, Src: 76:2d:1c:43:96:bd (76:2d:1c:43:96:bd), Dst: 56:44:d0:06:59:33 (56:44:d0:06:59:33)
> Internet Protocol Version 4, Src: 10.10.10.25, Dst: 10.10.10.26 Worker01 --> Worker02 물리 IP
User Datagram Protocol, Src Port: 48384, Dst Port: 4789
Source Port: 48384
Destination Port: 4789 **VXLAN Port (UDP)**
Length: 124
Checksum: 0xaebe5 [unverified]
[Checksum Status: Unverified]
[Stream index: 0]
> [Timestamps]
UDP payload (116 bytes)
Virtual eXtensible Local Area Network
Flags: 0x0000, VXLAN Network ID (VNI)
Group Policy ID: 0
VXLAN Network Identifier (VNI): 4096 **VNI 식별**
Reserved: 0
> Ethernet II, Src: 66:8c:33:86:44:cc (66:8c:33:86:44:cc), Dst: 66:f7:9a:22:22:c3 (66:f7:9a:22:22:c3)
> Internet Protocol Version 4, Src: 192.168.5.11, Dst: 192.168.49.2 **Calico VXLAN Interface**
User Datagram Protocol, Src Port: 47490, Dst Port: 53
Domain Name System (query)