

[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
```

```
calicoctl get ippool -o wide
```

NAME	CIDR	NAT	IPIPMode
VXLANMode	DISABLED	DISABLEBGPEXPORT	SELECTOR
default-ipv4-ippool	192.168.0.0/16	true	Always
false	false		Never

```
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.yml 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"
  .

#
kubectll 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 calicoctl node status
Calico process is running.
The BGP backend process (BIRD) is not running.
```

```
# Network VXLANMODE 가
calicoctl get ippool -o wide
NAME CIDR NAT IPIPMODE
VXLANMODE DISABLED DISABLEBGPEXPORT SELECTOR
default-ipv4-ippool 192.168.0.0/16 true Never
Always false false
all()
```

```
# tunl0 가 vxlan 가
# vxlan 가
```

```
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 cali1fdac863dc5 // Node(Worker02)가 GW
```

Worker

```
hostway@controller:~$ ip netns | 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	NOMINATED	NODE	READINESS GATES	
sampleos-646dc9654b-8xjw9	1/1	Running	0	45s
192.168.5.11	<none>	worker01	<none>	
sampleos-646dc9654b-gxn75	1/1	Running	0	45s
192.168.5.10	<none>	worker01	<none>	
sampleos-646dc9654b-snkxg	1/1	Running	0	45s
192.168.30.75	<none>	worker02	<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 .



[] CentOS 7 Kubernetes Install

CentOS 7 Kubernetes

OS : CentOS 7.6.1810 Minimal
Account : root
- SNAT IP
Controller : 10.10.10.237 SSH:4223
Worker-01 : 10.10.10.204 SSH:4224
Worker-02 : 10.10.10.190 SSH:4225

```
# root . sudo
useradd -d /home/username username
echo "password" | passwd username --stdin

# su
chmod 700 /usr/bin/su

# sudoer wheel 가
sed -ie '/wheel/s/$/\:username/' /etc/group

# Timezone
sudo timedatectl set-timezone Asia/Seoul

# SWAP OFF
sudo swapoff -a
sudo sed -i -e '/swap/d' /etc/fstab
```

```
# firewalld off
sudo systemctl stop firewalld && sudo systemctl disable
firewalld

# Selinux
setenforce 0
sudo sed -i 's/SELINUX=enforcing/SELINUX=disabled/g'
/etc/selinux/config

# Hostname
sudo hostnamectl set-hostname controller
sudo hostnamectl set-hostname worker-01
sudo hostnamectl set-hostname worker-02

## Controller / Worker
#curl -s https://get.docker.com | sudo sh
curl -fsSL https://get.docker.com -o get-docker.sh
sudo sh get-docker.sh

## Check
sudo docker -v
sudo docker ps -a

## Controller / Worker
sudo mkdir /etc/docker
cat <<EOF | sudo tee /etc/docker/daemon.json
{
  "exec-opts": ["native.cgroupdriver=systemd"],
  "log-driver": "json-file",
  "log-opts": {
    "max-size": "100m"
  },
  "storage-driver": "overlay2"
}
EOF

## Docker enable && restart
sudo systemctl enable docker
sudo systemctl daemon-reload
sudo systemctl restart docker
```

```
## Packages Repo
sudo cat <<EOF | sudo tee /etc/yum.repos.d/kubernetes.repo
[kubernetes]
name=Kubernetes
baseurl=https://packages.cloud.google.com/yum/repos/kubernetes
-el7-x86_64
enabled=1
gpgkey=https://packages.cloud.google.com/yum/doc/yum-key.gpg
https://packages.cloud.google.com/yum/doc/rpm-package-key.gpg
EOF
```

```
## Install
sudo yum install -y kubelet kubeadm kubectl --
disableexcludes=kubernetes
```

Controller Init

```
# Controller. IP API
(Advertise)
sudo kubeadm init --ignore-preflight-errors=all --pod-network-
cidr=192.168.0.0/16 --apiserver-advertise-address=10.10.10.237
```

```
# Regular User Privileges
mkdir -p $HOME/.kube
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config
```

```
# Network Plugin Setting ( Calico )
curl
https://projectcalico.docs.tigera.io/manifests/calico.yaml -O
kubectl apply -f calico.yaml
```

```
# System Namespace ( kube-system ) check. CoreDNS 가
kubectl get pods -o wide -A
```

NAMESPACE	NAME	READY		
STATUS	RESTARTS	AGE	IP	NODE
NOMINATED	NODE	READINESS	GATES	
kube-system	calico-kube-controllers-7c845d499-p85pm	1/1		
Running	0	3m6s	192.168.49.3	controller

```

<none>                <none>
kube-system           calico-node-fnm2q           1/1
Running 0             3m6s    10.10.10.237    controller
<none>                <none>
kube-system           coredns-64897985d-cgvml    1/1
Running 0             5m41s    192.168.49.2    controller
<none>                <none>
kube-system           coredns-64897985d-vdckf    1/1
Running 0             5m42s    192.168.49.1    controller
<none>                <none>
kube-system           etcd-controller            1/1
Running 0             5m54s    10.10.10.237    controller
<none>                <none>
kube-system           kube-apiserver-controller   1/1
Running 0             5m54s    10.10.10.237    controller
<none>                <none>
kube-system           kube-controller-manager-controller 1/1
Running 0             6m       10.10.10.237    controller
<none>                <none>
kube-system           kube-proxy-nn5zn           1/1
Running 0             5m42s    10.10.10.237    controller
<none>                <none>
kube-system           kube-scheduler-controller   1/1
Running 0             5m54s    10.10.10.237    controller
<none>                <none>

```

```

# ( ) Multi NIC 가 INTERNAL-IP
가 K8S NIC IP 가
INTERNAL-IP
INTERNAL-IP Init
kubeadm --apiserver-advertise-address IP

```

```

cat << EOF | sudo tee /etc/default/kubelet
KUBELET_EXTRA_ARGS='--node-ip $(hostname -I | cut -d ' ' -f2)'
EOF
sudo systemctl daemon-reload
sudo systemctl restart kubelet
kubectl cluster-info

```


Worker Join

```
# Worker-01 Woker-02 Node User Privileges

sudo scp /etc/kubernetes/admin.conf
username@10.10.10.204:/home/username/admin.conf
sudo scp /etc/kubernetes/admin.conf
username@10.10.10.190:/home/username/admin.conf

# Worker
mkdir -p $HOME/.kube
sudo cp -i ./admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config

# Worker kubeadm Join .
sudo kubeadm join 10.10.10.237:6443 --token
jgocer.fu65ql39kdod5qi0 \
--discovery-token-ca-cert-hash
sha256:3cb85267e89913d7865d219922daaa8fc6e788dd2be0e2f80fae271
76e2dfe3b

#
kubeadm token create --print-join-command

# Check
kubectl get nodes -o wide
NAME STATUS ROLES AGE VERSION
INTERNAL-IP EXTERNAL-IP OS-IMAGE KERNEL-
VERSION CONTAINER-RUNTIME
controller Ready control-plane,master 16m v1.23.5
10.10.10.237 <none> CentOS Linux 7 (Core)
3.10.0-1062.el7.x86_64 docker://20.10.14
worker-01 Ready <none> 55s v1.23.5
10.10.10.204 <none> CentOS Linux 7 (Core)
3.10.0-1062.el7.x86_64 docker://20.10.14
worker-02 NotReady <none> 38s v1.23.5
10.10.10.190 <none> CentOS Linux 7 (Core)
3.10.0-1062.el7.x86_64 docker://20.10.14

# Check Pod Create
kubectl run hello --image=nginx --dry-run=client -o yaml |
```

```

kubectly apply -f-
pod/hello created
[myungin.baek@controller ~]$ kubectly get pods -o wide
NAME          READY    STATUS    RESTARTS    AGE    IP
NODE          NOMINATED NODE    READINESS GATES
hello        1/1     Running    0           42s    192.168.171.1
worker-01    <none>          <none>

```

[OS] CentOS 7 iptables

iptables

CentOS 7 , SSH

(Pre) CentOS 7 firewalld iptables

```

firewalld , iptables iptables.target
service .

```

```

# firewalld disable
systemctl stop firewalld && systemctl disable firewalld

```

```

# firewalld service .
# /etc/sysconfig/iptables

```

```

yum install iptables-services
service iptables reload
service iptables status

```

```

#
service iptables save

```

```

#
service iptables reload

#
#           .           -c ( ALL Rule )
#           ROUTE(NAT)           .
iptables-save -c > rules.txt

#
iptables-restore < rules.txt

iptables           (           IP           )

#
iptables -F

# lo           ACCEPT
iptables -A INPUT -i lo -j ACCEPT

#           IP           .           (SSH)           -p tcp (-m
tcp           가 ) --dport 22           가
iptables -A INPUT -s 1.2.3.4/32 -m comment --comment "           " -j
ACCEPT

# state           ACCEPT.
iptables -A INPUT -m state --state RELATED,ESTABLISHED -j
ACCEPT

# (           ) Ping request           가           .           가
iptables -A INPUT -j REJECT --reject-with icmp-host-prohibited

# (           ) Ping request           .
iptables -A INPUT -p icmp --icmp-type echo-request -j REJECT

# (           ) Ping           DROP.           ACCEPT
iptables -A INPUT -p icmp -j DROP

#           TCP           DROP
iptables -A INPUT -p tcp -j DROP

```

```
#
service iptables save
```

가 가

```
# -A 가 DROP Line 가 Line
# -I INPUT [DROP Line] DROP 가 .
```

```
iptables -nL --line-number
```

```
-----
Chain INPUT (policy ACCEPT)
```

num	target	prot	opt	source	destination
1	ACCEPT	all	--	1.2.3.4	0.0.0.0/0
/* */					
2	DROP	tcp	--	0.0.0.0/0	0.0.0.0/0

```
# 2 DROP 가 .
iptables -I INPUT 2 -s 5.6.7.8 -j ACCEPT -m comment --comment "
가"
```

```
iptables -nL --line-number
```

```
-----
Chain INPUT (policy ACCEPT)
```

num	target	prot	opt	source	destination
1	ACCEPT	all	--	1.2.3.4	0.0.0.0/0
/* */					
2	ACCEPT	all	--	5.6.7.8	0.0.0.0/0
/* 가 */					
3	DROP	tcp	--	0.0.0.0/0	0.0.0.0/0

```
# /etc/sysconfig/iptables
reload 가 .
```

```
iptables -D INPUT [Number]
```

[] CNI - Calico Plugin

: CNI Calico Network

#1 (controller , worker)

CNI (Container Network Interface)

CNCF
Kubernetes Plugin Kubenet CNI Network

Calico Network?

vRouter (L3)
Kubernetes Plugin Network CNI Network

Document URL :
<https://projectcalico.docs.tigera.io/reference/>

Non-overlay

Direct
- BGP(Border Gateway Protocol) BIRD

```

Pod Node
Pod Calico Pod BGP Peer
가
. ( ex:
)

```

Overlay Network

```

Workload IP( ex:
)
(Encaptulation) (L2)

```

```

: Node IP 가 , POD
IP 가

```

```

# IP in IP (Default)
- 가 Direct
IP tunl0(tunneling)
가
Direct 가 BGP (BIRD)
Node (IPVS)
Calico Routing

```

```

# VXLAN
- 가

```

```

IP in IP
. ( ex: Azure )
Calico BGP 가
VXLAN Node
L2 UDP
IP in IP 가

```

```

# Cross-subnet
가 ( 가 ) 가
( / )
( )

```

```

# WireGuard

```

Calico

가

가

Calicoctl

```

Controller      Calico Network
Host            kubectl  plugin

```

```

# Host
$ cd /usr/local/bin
$ sudo curl -L https://github.com/projectcalico/calico/releases/download/v3.22.1/calicoctl-linux-amd64 -o calicoctl
$ sudo chmod +x calicoctl

```

Check

```

Calico 가      Network Pool Block
$ sudo calicoctl ipam show --show-blocks
+-----+-----+-----+-----+-----+
| GROUPING |          CIDR          | IPS TOTAL | IPS IN USE |
IPS FREE  |
+-----+-----+-----+-----+-----+
| IP Pool  | 192.168.0.0/16        | 65536 | 5 (0%)    |
65531 (100%) |
| Block   | 192.168.136.0/26     | 64 | 4 (6%)   | 60
(94%)      |
| Block   | 192.168.153.192/26   | 64 | 1 (2%)   | 63
(98%)      |
+-----+-----+-----+-----+-----+

```

BGP

```

$ sudo calicoctl node status
Calico process is running.
IPv4 BGP status

```

```

+-----+-----+-----+-----+-----+
| PEER ADDRESS | PEER TYPE | STATE | SINCE |
INFO      |

```

```

+-----+-----+-----+-----+-----+
-----+
| 203.248.23.215 | node-to-node mesh | up      | 05:27:05 |
Established |
+-----+-----+-----+-----+-----+
-----+

```

Block

```
$ route -n | egrep "tun|cali|\\"*
```

```

192.168.136.0    0.0.0.0          255.255.255.192 U        0
0              0 *
192.168.136.1   0.0.0.0          255.255.255.255 UH       0        0
0 calibc6c3028870
192.168.136.2   0.0.0.0          255.255.255.255 UH       0        0
0 calid6edae09645
192.168.136.3   0.0.0.0          255.255.255.255 UH       0        0
0 calic6bfd11bfbe
192.168.153.192 203.248.23.215  255.255.255.192 UG       0        0
0 tunl0

```

Pod가 calicxxxxxx

System(default) Namespace -A 가 .

```
$ calicoctl get workloadendpoint -A
```

NAMESPACE	WORKLOAD	INTERFACE	NETWORKS	NODE
kube-system	calico-kube-controllers-56fcbf9d6b-nlqg2	calid6edae09645	192.168.136.2/32	user-
kube-system	coredns-64897985d-jgj5s	calic6bfd11bfbe	192.168.136.3/32	user-
kube-system	coredns-64897985d-vbpn4	calibc6c3028870	192.168.136.1/32	user-

Calico Veth type(Pair) .

```

$ ip -br -c link show type veth
calibc6c3028870@if3 UP ee:ee:ee:ee:ee:ee
<BROADCAST,MULTICAST,UP,LOWER_UP>
calid6edae09645@if4 UP ee:ee:ee:ee:ee:ee
<BROADCAST,MULTICAST,UP,LOWER_UP>
calic6bfd11bfbe@if4 UP ee:ee:ee:ee:ee:ee
<BROADCAST,MULTICAST,UP,LOWER_UP>

```


Calico Management Pod

Daemon Pod
Controller Worker Node Pod 가

```
$ kubectl get pods -o wide -n kube-system
NAME                                READY   STATUS
RESTARTS   AGE    IP                NODE
calico-kube-controllers-56fcbf9d6b-nlqg2  1/1     Running    0
30m    192.168.136.2    user-controller
calico-node-8cts6                          1/1     Running    0
30m    10.0.2.15        user-controller
calico-node-mb9n6                          1/1     Running    0
29m    10.0.2.15        user-worker
```

Calico DB etcd datastore

```
$ kubectl get pods -o wide -n kube-system | grep -i etcd
etcd-user-controller                1/1     Running    0
39m    10.0.2.15        user-controller
```

Calico Felix

Pod kube-proxy
etcd Pod Network
kube-proxy 가 iptables / ipvs Mode
iptables ipvs

□ IPVS = Hash

```
$ sudo iptables -t nat -S | grep -i cali
$ sudo iptables -t filter -S | grep -i cali
```

Networking

IP in IP Networking

Controller Node Worker Node Pod

❌
1) Controller 192.168.136.2 Pod Worker 192.168.153.193
Pod

```

2) Controller Pod (veth) Pair Host calico
(veth) ARP
3) Host Calico Worker Pod
ARP
4) Calico link-local ( )
HOST 가 BIRD Worker
5) Controller Calico vRouter ARP_Proxy
Worker ARP
6) BIRD Tunl0 --> Host Pod
가
7)

Felix SNAT ( MASQUERADE ) tunl0
HOST ens33 .

```

Packet Check

```

# ( Controllor POD <---> Worker POD ) Ping

$ kubectl get pod -o wide
NAME READY STATUS RESTARTS AGE IP
NODE NOMINATED NODE READINESS GATES
hello-776c774f98-894tt 1/1 Running 0 13d
192.168.153.193 user-worker <none> <none>
hi 1/1 Running 0 13d
192.168.136.5 user-controller <none> <none>

# Worker POD --> Container POD. Ping Pod
Host PID .
$ sudo nsenter -t 225201 -n ping 192.168.136.5
64 bytes from 192.168.136.5: icmp_seq=627 ttl=62 time=0.709 ms
64 bytes from 192.168.136.5: icmp_seq=628 ttl=62 time=0.675 ms
64 bytes from 192.168.136.5: icmp_seq=629 ttl=62 time=0.727 ms
64 bytes from 192.168.136.5: icmp_seq=630 ttl=62 time=0.797 ms
64 bytes from 192.168.136.5: icmp_seq=631 ttl=62 time=0.887 ms

# Controller . IPIP API
, API .
$ sudo tcpdump -i enp0s8 -nn proto 4 -w test.pcap

# Wireshark .

```

1) POD IP ICMP



2) MAC Controller Worker Node API IP

Controller

```
3: enp0s8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether 08:00:27:39:ce:bd brd ff:ff:ff:ff:ff:ff
    inet 203.248.23.214/25 brd 203.248.23.255 scope global enp0s8
```

Worker

```
3: enp0s8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether 08:00:27:bc:85:3a brd ff:ff:ff:ff:ff:ff
    inet 203.248.23.215/25 brd 203.248.23.255 scope global enp0s8
```

3) IPv4 Protocol 2

```
# Outer IP POD Inner IP 2
```



Outer IP 가 InnerIP IP-IP Protocol



4) Messages 가



Network Overlay : <https://ikcoo.tistory.com/117>

CentOS 7

Windows RDP

CentOS 7 Windows RDP

```
# OS
CentOS 7.9 x86_64 minimal
--> XRDP          GUI    -->    GUI    Windows
RDP
```

Linux GUI

```
# GUI GroupInstall
root@localhost ~]# yum groups list | grep -i desktop
  Cinnamon Desktop
  MATE Desktop
  GNOME Desktop
  General Purpose Desktop
  LXQt Desktop
# GNOME          "Server with GUI"
root@localhost ~] yum groupinstall "GNOME Desktop"

# GUI init
[root@localhost ~]# systemctl get-default
multi-user.target
[root@localhost ~]# systemctl set-default graphical.target
[root@localhost ~]# systemctl get-default
graphical.target
# Reboot          GUI
[root@localhost ~]# reboot
```

Linux

```
# XRDP Install.
```

```
[root@localhost ~]# yum install epel-release
[root@localhost ~]# yum install xrdp
[root@localhost ~]# systemctl enable xrdp && systemctl start
xrdp
```

```
# selinux disable iptables -F or tcp/3389 가
```

rdesktop

```
# openssl-devel
yum -y install gcc openssl-devel
wget
https://github.com/rdesktop/rdesktop/releases/download/v1.8.6/
rdesktop-1.8.6.tar.gz
tar xvzf rdesktop-1.8.6.tar.gz
cd rdesktop-1.8.6/
./configure --disable-credssp --disable-smartcard
make
make install
```

Check

```
# RDP , rdesktop -u [User] [ip]
root@localhost ~]# rdesktop -u administrator 10.10.10.5
Autoselected keyboard map en-us
Connection established using SSL.
WARNING: Remote desktop does not support colour depth 24;
falling back to 16
```

