□□□□ □□ □ □□ (Kubeadm)

□□□□ **Vanilla Kubernetes**"□ □□□□□ □□ □□□ □□□ □□□□□□□ □□ □□□ □□□□□. □ □□□ □□□□□□□ □□ □□□□□ □□ □□□□□, □□ □□□□□□ □ □□ □□□

□□□ *VMWare Workstation*□□ □□ □□

- Master Node: 192.168.1.10
- Worker Node 1: 192.168.1.20
- Worker Node 2: 192.168.1.30

> □□□□□ □□□□ □□ □ Swap Memory □□□□□ □□ □□. Swap Memory□ □□ □□□□ □□□□ RAM□ □□□□ □□ □□□□□□, K8s□ □
> □□□□ □□□□□□ □□. □ □□□ □□□ □□ □□□

CentOS 7□□ K8s □□□□

Docker Installation

□□ □, □□ □□□□ □□ □□

```
sudo dnf remove docker \
                  docker-client \
                  docker-client-latest \
                  docker-common \
                  docker-latest \
                  docker-latest-logrotate \
                  docker-logrotate \
                  docker-engine
```

yum-utils □□□ □□ yum repository (□□□) □□

```
sudo dnf install -y yum-utils
```

```
sudo dnf-config-manager \
    --add-repo \
    https://download.docker.com/linux/centos/docker-ce.repo
```

□□ □□ □□ □□

```
sudo dnf install docker-ce docker-ce-cli containerd.io docker-buildx-plugin docker-compose-plugin
```

□□ □□□ □□

```
sudo systemctl start docker
```

□□ □□ □□

```
docker version
```

Kubernetes Installation

Master Node□ *Worker Node*□□ □□ □□ □□

Swap Memory □□□□

```
swapoff -a
```

daemon.json □□ □□ □ cgroupdriver□ systemd□ □□

```
vi /etc/docker/daemon.json


{
    "exec-opts": ["native.cgroupdriver=systemd"]
}
```

daemon □□□

```
systemctl daemon-reload
systemctl restart docker
```

□□□□□ □□ □□□ □□ □ □□

```
vi /etc/yum.repos.d/kubernetes.repo

[kubernetes]
name=Kubernetes
baseurl=https://packages.cloud.google.com/yum/repos/kubernetes-el7-x86_64
enabled=1
gpgcheck=0
repo_gpgcheck=0
gpgkey=https://packages.cloud.google.com/yum/doc/yum-key.gpg
       https://packages.cloud.google.com/yum/doc/rpm-package-key.gpg

# □□□□□ □□
dnf install -y kubelet-1.19.16-0.x86_64 kubectl-1.19.16-0.x86_64 kubeadm-1.19.16-0.x86_64
```

□□□□□ □□ □□ □□

```
rpm -qa | grep kube
```

□□□□□ □□□□ □□□ □□□ □□ □□□ □□ □□

```
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=443/tcp
firewall-cmd --permanent --add-port=2376/tcp
firewall-cmd --permanent --add-port=2379/tcp
firewall-cmd --permanent --add-port=2380/tcp
firewall-cmd --permanent --add-port=6443/tcp
firewall-cmd --permanent --add-port=8472/udp
```

```
firewall-cmd --permanent --add-port=9099/tcp
firewall-cmd --permanent --add-port=10250/tcp
firewall-cmd --permanent --add-port=10251/tcp
firewall-cmd --permanent --add-port=10252/tcp
firewall-cmd --permanent --add-port=10254/tcp
firewall-cmd --permanent --add-port=10255/tcp
firewall-cmd --permanent --add-port=30000-32767/tcp
firewall-cmd --permanent --add-port=30000-32767/udp
firewall-cmd --permanent --add-masquerade
firewall-cmd --reload
```

□□□ □□□□□ □□

□□□□□ □□□□ □□

```
kubeadm init --apiserver-advertise-address=192.168.1.10 --pod-network-cidr=10.244.0.0/16

# K8s control plane □□ □□ □□
Your Kubernetes control-plane has initialized successfully!

To start using your cluster, you need to run the following as a regular user:

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

Alternatively, if you are the root user, you can run:

  export KUBECONFIG=/etc/kubernetes/admin.conf

You should now deploy a pod network to the cluster.
Run "kubectl apply -f [podnetwork].yaml" with one of the options listed at:
  https://kubernetes.io/docs/concepts/cluster-administration/addons/

Then you can join any number of worker nodes by running the following on each as root:

kubeadm join 192.168.1.10:6443 --token 172vji.r0u77jcmcnccm6no \
```

```
    --discovery-token-ca-cert-hash
sha256:72b9648c647f724ab52471847cb06c47b23097375f2e67633b745fc69db16e8d
```

kubectl □□□ □□ admin.conf □□ □□

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

□□□ Flannel □□□ (

[CNI(Container Network Interface)□ □□□□□ □□ □□□□□ □□□ □ □□ □□□□□ □□□ □□ □□□□□ □□. □□□□ □□ □□ CNI□□□□□ □□

□ □□□ □□□

```
curl -O -L https://raw.githubusercontent.com/coreos/flannel/master/Documentation/kube-flannel.yml
```

kube-flannel.yml □□ □□ --iface=(□□ □□□□□ □□) □□

```
vi kube-flannel.yml

args:
        - --ip-masq
        - --kube-subnet-mgr
        - --iface=ens160
```

□□□□ □□□□ YAML □□ □□ □ kubelet □□□

```
kubectl apply -f kube-flannel.yml
systemctl restart kubelet
```

□□*Worker Node*□□ □□

Master Node□□ □□ Discovery Token □□ join □□□□ □□ Master□ Worker □□ □□

```
kubeadm join 192.168.1.10:6443 --token 172vji.r0u77jcmcnccm6no \
    --discovery-token-ca-cert-hash
sha256:72b9648c647f724ab52471847cb06c47b23097375f2e67633b745fc69db16e8d

This node has joined the cluster:
* Certificate signing request was sent to apiserver and a response was received.
* The Kubelet was informed of the new secure connection details.

Run 'kubectl get nodes' on the control-plane to see this node join the cluster.
```

Join □□□ □□ □ □□□ □□□□ □□□□□ □□□□ □□

```
kubectl get nodes
NAME     STATUS   ROLES    AGE    VERSION
master   Ready    master   107m   v1.19.16
node-1   Ready    <none>   91s    v1.19.16
node-2   Ready    <none>   48s    v1.19.16
```

Rocky Linux 9.1□□ K8s □□□□

Docker Installation

□□□□ □□ □ □□□□ □□□ □□

```
sudo dnf check-update
sudo dnf update
```

□□ repository (□□□) □□

```
sudo dnf config-manager --add-repo https://download.docker.com/linux/centos/docker-ce.repo
```

□□ □□

```
sudo dnf install docker-ce docker-ce-cli containerd.io
```

□□ □□ □ □□

```
sudo systemctl start docker
sudo systemctl enable docker

sudo docker version
```

'sudo' □□ □□ docker □ □□ □□□

```
sudo usermod -aG docker $(whoami)
```

Kubernetes Installation

□□*Master Node*□ *Worker Node*□□ □□ □□ □□

Swap Memory □□□□

```
sudo swapoff -a
```

daemon.json □□ □□ □ cgroupdriver□ systemd□ □□

```
sudo vi /etc/docker/daemon.json

{
   "exec-opts": ["native.cgroupdriver=systemd"]
}
```

daemon □□□

```
sudo systemctl daemon-reload
sudo systemctl restart docker
```

□□□□□ □□ □□□ □□ □ □□

```
sudo vi /etc/yum.repos.d/kubernetes.repo

[kubernetes]
name=Kubernetes
baseurl=https://packages.cloud.google.com/yum/repos/kubernetes-el7-x86_64
enabled=1
gpgcheck=0
repo_gpgcheck=0
gpgkey=https://packages.cloud.google.com/yum/doc/yum-key.gpg
       https://packages.cloud.google.com/yum/doc/rpm-package-key.gpg

# □□□□□ □□
sudo dnf -y kubelet-1.19.16-0.x86_64 kubectl-1.19.16-0.x86_64 kubeadm-1.19.16-0.x86_64
```

□□□□□ □□ □□ □□

```
sudo rpm -qa | grep kube
```

□□□□□ □□□□ □□□ □□□ □□ □□□ □□ □□

```
sudo firewall-cmd --permanent --add-port=80/tcp
sudo firewall-cmd --permanent --add-port=443/tcp
sudo firewall-cmd --permanent --add-port=2376/tcp
sudo firewall-cmd --permanent --add-port=2379/tcp
sudo firewall-cmd --permanent --add-port=2380/tcp
sudo firewall-cmd --permanent --add-port=6443/tcp
sudo firewall-cmd --permanent --add-port=8472/udp
```

```
sudo firewall-cmd --permanent --add-port=9099/tcp
sudo firewall-cmd --permanent --add-port=10250/tcp
sudo firewall-cmd --permanent --add-port=10251/tcp
sudo firewall-cmd --permanent --add-port=10252/tcp
sudo firewall-cmd --permanent --add-port=10254/tcp
sudo firewall-cmd --permanent --add-port=10255/tcp
sudo firewall-cmd --permanent --add-port=30000-32767/tcp
sudo firewall-cmd --permanent --add-port=30000-32767/udp
sudo firewall-cmd --permanent --add-masquerade
sudo firewall-cmd --reload
```

□□□ □□□□□ □□

□□ □□ (root □□)□□ □□

```
sudo su
```

□□□□□ □□□□ □□

```
kubeadm init --apiserver-advertise-address=192.168.1.10 --pod-network-cidr=10.244.0.0/16

# K8s control plane □□ □□ □□
Your Kubernetes control-plane has initialized successfully!

To start using your cluster, you need to run the following as a regular user:

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

Alternatively, if you are the root user, you can run:

  export KUBECONFIG=/etc/kubernetes/admin.conf

You should now deploy a pod network to the cluster.
Run "kubectl apply -f [podnetwork].yaml" with one of the options listed at:
```

```
  https://kubernetes.io/docs/concepts/cluster-administration/addons/

Then you can join any number of worker nodes by running the following on each as root:

sudo kubeadm join 192.168.1.10:6443 --token 1x7qb2.dww3u7mjsrq2hoxt --discovery-token-ca-cert-hash
sha256:10a6803e9a45bb029af4ad3c1d0d894dfaee9980d2318c495739296daeffb9eb
```

kubectl □□□□ □□ admin.conf □□ □□ (□ root□□□□ □□□!)

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

□□□□ □□□□ □□□□

```
curl -O -L https://raw.githubusercontent.com/coreos/flannel/master/Documentation/kube-flannel.yml
```

□□□□ □□□□ YAML □□ □□ □ kubelet □□□ (□ root□□□□ □□□!)

```
kubectl apply -f kube-flannel.yml
systemctl restart kubelet
```

□□*Worker Node*□□ □□

Master Node□□ □□ Discovery Token □□ join □□□□ □□ Master□ Worker □□ □□

```
sudo kubeadm join 192.168.1.10:6443 --token 1x7qb2.dww3u7mjsrq2hoxt \
--discovery-token-ca-cert-hash
sha256:10a6803e9a45bb029af4ad3c1d0d894dfaee9980d2318c495739296daeffb9eb

This node has joined the cluster:
* Certificate signing request was sent to apiserver and a response was received.
* The Kubelet was informed of the new secure connection details.
```

```
Run 'kubectl get nodes' on the control-plane to see this node join the cluster.
```

Join □□□ □□ □ □□□ □□□□ □□□□□ □□□□ □□

```
sudo kubectl get nodes
NAME      STATUS   ROLES    AGE    VERSION
master    Ready    master   33m    v1.19.16
worker-1  Ready    <none>   91s    v1.19.16
worker-2  Ready    <none>   48s    v1.19.16
```

Ubuntu 20.04.6□□ K8s □□□□

Docker Installation

□□□□ □□ □ □□

```
sudo apt update
```

□□ □□□ □□

```
sudo apt install ca-certificates curl gnupg lsb-release
```

□□ GPG □ □□

```
curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo gpg --dearmor -o /usr/share/keyrings/docker-archive-keyring.gpg
```

□□ Repository □□

```
echo \
  "deb [arch=$(dpkg --print-architecture) signed-by=/usr/share/keyrings/docker-archive-keyring.gpg]
https://download.docker.com/linux/ubuntu \
  $(lsb_release -cs) stable" | sudo tee /etc/apt/sources.list.d/docker.list > /dev/null
```

□□ □□ □□ □ □□

```
sudo apt update
sudo apt install docker-ce docker-ce-cli containerd.io

sudo docker version
sudo systemctl status docker

## □□ □□□ □ docker □ □□□□ □□□□□ □□ ##
sudo systemctl enable docker
```

Kubernetes Installation

Master Node□ *Worker Node*□□ □□ □□ □□

Swap Memory □□□□

```
sudo swapoff -a
```

□□□ □□□ □□ iptables□ □□□ □□ □□ □□

```
sudo vi /etc/modules-load.d/k8s.conf

br_netfilter

sudo vi /etc/sysctl.d/k8s.conf

net.bridge.bridge-nf-call-ip6tables = 1
net.bridge.bridge-nf-call-iptables = 1
```

```
# sysctl.conf □□ □ □□□□ □□ □□ □ □ □□
sudo vi /etc/sysctl.conf

net.bridge.bridge-nf-call-ip6tables = 1
net.bridge.bridge-nf-call-iptables = 1

# sysctl □□ □□□□ □□□ □□ □□ □□
sudo sysctl --system
sudo sysctl -p
```

daemon.json □□ □□ □ cgroupdriver□ systemd□ □□

```
sudo vi /etc/docker/daemon.json


{
  "exec-opts": ["native.cgroupdriver=systemd"],
  "log-driver": "json-file",
  "log-opts": {
    "max-size": "100m"
  },
  "storage-driver": "overlay2"
}
```

daemon □□□

```
sudo systemctl daemon-reload
sudo systemctl restart docker
sudo systemctl enable docker
```

apt □□□□ □ ca □□ □□□ □□□□ □ □

```
sudo apt update
sudo apt install -y apt-transport-https ca-certificates curl
sudo curl -fsSLo /usr/share/keyrings/kubernetes-archive-keyring.gpg
https://packages.cloud.google.com/apt/doc/apt-key.gpg
```

□□ gpg □ □□ □□□ □□ □□ □ apt □□□

```
echo "deb [signed-by=/usr/share/keyrings/kubernetes-archive-keyring.gpg] https://apt.kubernetes.io/
kubernetes-xenial main" | sudo tee /etc/apt/sources.list.d/kubernetes.list
sudo apt update
```

kubeadm, kubelet □□□ kubectl □□ □ □□□

```
sudo apt install -y kubelet kubeadm kubectl
sudo apt-mark hold kubelet kubeadm kubectl

sudo systemctl start kubelet && systemctl enable kubelet
```

□□□ □□□□□ □□

□□□□□ □□□□ □□□ □□ kubeadm init □□

```
sudo kubeadm init --apiserver-advertise-address=192.168.1.10 --pod-network-cidr=10.244.0.0/16

# K8s control plane □□ □□ □□
Your Kubernetes control-plane has initialized successfully!

To start using your cluster, you need to run the following as a regular user:

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

Alternatively, if you are the root user, you can run:

  export KUBECONFIG=/etc/kubernetes/admin.conf

You should now deploy a pod network to the cluster.
Run "kubectl apply -f [podnetwork].yaml" with one of the options listed at:
```

```
  https://kubernetes.io/docs/concepts/cluster-administration/addons/

Then you can join any number of worker nodes by running the following on each as root:

kubeadm join 192.168.1.10:6443 --token mu72kx.sn06xzcg3lde7mha \
--discovery-token-ca-cert-hash
sha256:d134654458b474c796667776ab8175adbc0e1dc4bc21a712a19a7bb405ee9273
```

kubectl □□□□ □□ admin.conf □□ □□

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

□□□□ □□□□ □□□□

```
curl -O -L https://raw.githubusercontent.com/coreos/flannel/master/Documentation/kube-flannel.yml
```

□□□□ □□□□ YAML □□ □□ □ kubelet □□□

```
kubectl apply -f kube-flannel.yml
sudo systemctl restart kubelet
```

Worker Node□□ □□

Master Node□□ □□ Discovery Token □□ join □□□□ □□ Master□ Worker □□ □□

```
sudo kubeadm join 192.168.1.10:6443 --token mu72kx.sn06xzcg3lde7mha \
--discovery-token-ca-cert-hash
sha256:d134654458b474c796667776ab8175adbc0e1dc4bc21a712a19a7bb405ee9273

This node has joined the cluster:
* Certificate signing request was sent to apiserver and a response was received.
```

```
* The Kubelet was informed of the new secure connection details.

Run 'kubectl get nodes' on the control-plane to see this node join the cluster.
```

Join □□□ □□ □ □□□ □□□□ □□□□□ □□□□ □□

```
kubectl get nodes
NAME      STATUS   ROLES    AGE    VERSION
master    Ready    master   33m    v1.19.16
worker-1  Ready    <none>   91s    v1.19.16
worker-2  Ready    <none>   48s    v1.19.16
```

Error Troubleshooting

[ERROR CRI]: container runtime is not running [Issue Encountered]

□ □□□ □□□□ CRI□ Containerd□ □□ □□□ □□□□ kubeadm init □□□ □□□ □ □□ □□□□. □□□□ □□ config.toml □□□ □□□ □□.

□□

```
# config.toml □□□ □□ □□
sudo rm /etc/containerd/config.toml

# containerd □□□ □□□
sudo systemctl restart containerd
```

The connection to the server <IP □□>:6443 was refused - did you specif the right host or port?

□ □□□ □□ API □□□ □□ □□□ □□□ □□□□ □□□ □□ □□ □□□ □ □□. □□□ □ □□□ □□□ □□ □□□ □ □□ □□□ Swap □□□ □□□ □□□□ □□. □□ □□

□□

/etc/fstab □□ □□□□□ swap□ □□ □□ □□□

```
sudo vi /etc/fstab


# /etc/fstab: static file system information.
#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# <file system> <mount point>   <type>  <options>       <dump>  <pass>
# / was on /dev/sda3 during curtin installation
/dev/disk/by-uuid/05bb7e29-bd2d-4ffb-86c6-8868e48548f4 / ext4 defaults 0 1
# /boot/efi was on /dev/sda2 during curtin installation
/dev/disk/by-uuid/ebbd40d6-5f57-4214-806f-8cf8e929b23d /boot/efi ext4 defaults 0 1
# /swap.img     none    swap    sw      0       0
```

Revision #9
Created 29 May 2023 05:19:17 by □□
Updated 31 May 2023 00:46:39 by □□