

k3s

k3s rancher k8s k8s

<https://github.com/rancher/k3s/releases>

```
# 80, 443, 8080 6443

# bin
cp k3s /usr/local/bin/
# k3s 6443
# - docker k3s containerd
# --docker --no-deploy traefik traefik
# --no-deploy servicelb SLB
# --node-taint k3s-controlplane=true:NoExecute
# --disable-agent
# --no-deploy local-storage
./k3s server --docker
# token
cd /var/lib/rancher/k3s/server
cat node-token
cat /var/lib/rancher/k3s/server/node-token

# token
#
# --docker docker k3s containerd
k3s agent --token
K10aac59fd99494114e01b73264456d0be6093b8c908bcb36afe54b1da9aaa3b92f: : server: d05eb9b498f5da861e
ba8065d424a116 --server https://192.168.52.190:6443 --docker
# master
kubectl taint nodes 10.vpclub.io node-role.kubernetes.io/master=:NoExecute
#
kubectl taint nodes 10.vpclub.io node-role.kubernetes.io/master=:NoSchedule
```



```
# [ ] master [ ]

# vi /usr/lib/systemd/system/k3s.service
[Unit]
Description=Lightweight Kubernetes
Documentation=https://k3s.io
After=network-online.target
Wants=network-online.target
Requires=docker.socket

[Service]
Type=notify
# the default is not to use systemd for cgroups because the delegate issues still
# exists and systemd currently does not support the cgroup feature set required
# for containers run by docker
# [ ]master
#ExecStart=/usr/local/bin/k3s server --docker --node-taint k3s-controlplane=true:NoExecute
# [ ]worker
ExecStart=/usr/local/bin/k3s agent --docker --token
K10aac59fd99494114e01b73264456d0be6093b8c908bcb36afe54b1da9aaa3b92f::server:d05eb9b498f5da861e
ba8065d424a116 --server https://192.168.52.190:6443
ExecReload=/bin/kill -s HUP $MAINPID
TimeoutSec=0
RestartSec=2
Restart=always

# Note that StartLimit* options were moved from "Service" to "Unit" in systemd 229.
# Both the old, and new location are accepted by systemd 229 and up, so using the old location
# to make them work for either version of systemd.
StartLimitBurst=3

# Note that StartLimitInterval was renamed to StartLimitIntervalSec in systemd 230.
# Both the old, and new name are accepted by systemd 230 and up, so using the old name to make
# this option work for either version of systemd.
StartLimitInterval=60s

# Having non-zero Limit*s causes performance problems due to accounting overhead
```


- [redacted]

```
cat /etc/rancher/node/password
cat /var/lib/rancher/k3s/server/cred/node-passwd
```

```
# [redacted]
```

- [redacted]HTTPS

```
vi /var/lib/rancher/k3s/server/manifests/traefik.yaml
# [redacted]
ssl.enforced: "true"
# [redacted] ingress
ingress.kubernetes.io/ssl-redirect: "true"
# or
traefik.ingress.kubernetes.io/redirect-entry-point: https
```

```
# curl 80 [redacted]443 [redacted]
```

```
# [redacted] traefik [redacted]master
```

```
# [redacted] [redacted]pod[redacted]
```

```
# [redacted]tx[redacted] [redacted]k
ethtool -k flannel.1|grep tx-checksum-ip-generic
```

```
# [redacted]k
ethtool -K flannel.1 tx-checksum-ip-generic off
```

[redacted] #3

[redacted] [redacted] 28 [redacted] 2025 02:34:16

[redacted] [redacted] 5 [redacted] 2025 02:25:03