k3s

k3s是rancher精简的k8s实现，去除了大部分人用不到的功能，将一些组件进行集成，可以快速的搭建k8s集群

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

端口确认，主机的80,443,8080，6443端口不能被占用

复制到bin目录
cp k3s /usr/local/bin/
运行k3s，确保6443端口不被占用
增加--docker，k3s默认使用 containerd
--docker --no-deploy traefik 不部署traefik
--no-deploy servicelb 不部署内置SLB
--node-taint k3s-controlplane=true:NoExecute 服务器节点不可调度 # --disable-agent 默认情况下服务会成为一个代理
--no-deploy local-storage 本机存储驱动
. /k3s server --docker
查看toke
cd /var/lib/rancher/k3s/server
cat node-token
cat /var/lib/rancher/k3s/server/node-token

复制token
加入到集群
--docker 使用docker作为容器，k3s默认使用containerd
k3s agent --token
K10aac59fd99494114e01b73264456d0be6093b8c908bcb36afe54b1da9aaa3b92f::server:d05eb9b498f5da861e ba8065d424a116 --server https://192.168.52.190:6443 --docker
master 节点不参与调度并立即驱离已经调度的
kubectl taint nodes 10.vpclub.io node-role.kubernetes.io/master=:NoExecute
节点不参与调度，已经被调度的不受影响
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
```

```
# in the kernel.  We recommend using cgroups to do container-local accounting.
LimitNOFILE=infinity
LimitNPROC=infinity
LimitCORE=infinity

# Comment TasksMax if your systemd version does not support it.
# Only systemd 226 and above support this option.
TasksMax=infinity

# set delegate yes so that systemd does not reset the cgroups of docker containers
Delegate=yes

# kill only the docker process, not all processes in the cgroup
KillMode=process

[Install]
WantedBy=multi-user.target
```

□□kubectl□□

```
# □□k3s kubectl □□
vi ~/.bashrc
# □□
# alias kubectl='k3s kubectl'
# □□□□□
source ~/.bashrc
```

□□□

```
# □□□□□□
k3s server --docker --datastore-endpoint "mysql://root:123456@tcp(192.168.52.190:3307)/k3s"

#□□nginx□□□□□□□□□□□master□□□□□
```

- □□□□□□□□□□□□□□□□□□□□□□□□□□□□

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

# □□□□□□□□□□□□
```

- □□□□□□HTTPS

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



# curl 80  □443 □□□□□□□□□□

# □ traefik □□□master

# □□□□ □□□□□□□□pod□□

# □□tx□□ □□□□□□□□□k
ethtool -k  flannel.1| grep tx-checksum-ip-generic

# □□□□□□□□k
ethtool -K flannel.1 tx-checksum-ip-generic off
```

□□□ #3
□ □□□ □□□ 28 □□ 2025 02:34:16
□ □□□ □□□ 5 □□□ 2025 02:25:03