

K8s 故障：Pod 内无法访问外网

AUTHOR: 彭玲 TIME: 2022/3/14

K8s 故障：Pod 内无法访问外网

问题描述

原因分析

解决方案

南昌县环境部署问题解决

问题描述

在容器中对公网 (如 `baidu.com` 或 公网 IP) 执行 `ping` 操作, 无法 `ping` 通。

原因分析

`FLANNEL_SUBNET` 指定的子网段不在 `FLANNEL_NETWORK` 指定的大网段内。

```
fs@k8sworker08:~$ cat /run/flannel/subnet.env
FLANNEL_NETWORK=10.244.0.0/16
FLANNEL_SUBNET=10.8.11.1/24
FLANNEL_MTU=1450
FLANNEL_IPMASQ=true
```

通过 `kubectl -n kube-system edit configmap kubeadm-config` 查看 `kubeadm` 配置中 `podSubnet` 网络配置:

`podSubnet` 配置项应该与 `/run/flannel/subnet.env` 中的 `FLANNEL_NETWORK` 相同。

```

#
apiVersion: v1
data:
  ClusterConfiguration: |
    apiServer:
      extraArgs:
        authorization-mode: Node,RBAC
        timeoutForControlPlane: 4m0s
    apiVersion: kubeadm.k8s.io/v1beta2
    certificatesDir: /etc/kubernetes/pki
    clusterName: kubernetes
    controlPlaneEndpoint: 192.168.72.100:16443
    controllerManager: {}
    dns:
      type: CoreDNS
    etcd:
      local:
        dataDir: /var/lib/etcd
    imageRepository: registry.cn-hangzhou.aliyuncs.com/google_containers
    kind: ClusterConfiguration
    kubernetesVersion: v1.18.8
    networking:
      dnsDomain: cluster.local
      podSubnet: 10.8.0.0/16
      serviceSubnet: 10.244.0.0/16
    scheduler: {}
  ClusterStatus: |
    apiEndpoints:
      master01:
        advertiseAddress: 192.168.72.136
        bindPort: 6443
    apiVersion: kubeadm.k8s.io/v1beta2

```

解决方案

尝试重启 flannel 网络插件。

```

1 # 进入容器，正常情况下 ping 外网
2 anxin@node38:~$ kubectl exec -it emqx-deployment-6fdbf979b4-48nrc -n lk --
  sh
3 /opt/emqx $ sudo ping -c 3 baidu.com # ping 百度
4 PING baidu.com (220.181.38.148): 56 data bytes
5 64 bytes from 220.181.38.148: seq=0 ttl=51 time=33.863 ms
6 64 bytes from 220.181.38.148: seq=1 ttl=51 time=34.268 ms
7 64 bytes from 220.181.38.148: seq=2 ttl=51 time=33.587 ms
8
9 --- baidu.com ping statistics ---
10 3 packets transmitted, 3 packets received, 0% packet loss
11 round-trip min/avg/max = 33.587/33.906/34.268 ms

```

重启 flannel 网络插件的方式：

删除 flannel 对应的 Pod，利用 flannel 的 DaemonSet 自动启动 flannel 的 Pod。

```

1 # 删除 flannel 对应的 Pod
2 $ kubectl delete pods -n=kube-system --selector=app=flannel

```

确认 flannel 对应的 Pod 已自动启动，并正常运行：

```

1 # 查看 flannel 的 DaemonSet
2 anxin@node38:~$ kubectl get DaemonSet/kube-flannel-ds -n kube-system
3 NAME                DESIRED   CURRENT   READY   UP-TO-DATE   AVAILABLE   NODE
4 SELECTOR            AGE
5 kube-flannel-ds     3         3         3       3             3           <none>    63d
6
7 # 查看 flannel 对应的 Pod
8 anxin@node38:~$ kubectl get po -n kube-system --selector app=flannel -o wide
9 NAME                READY     STATUS    RESTARTS   AGE       IP
10 NODE                ...
11 kube-flannel-ds-7ctrb 1/1      Running   0          63d      10.8.30.35
12 node35
13 kube-flannel-ds-grvbw 1/1      Running   0          63d      10.8.30.38
14 node38
15 kube-flannel-ds-jxwxf 1/1      Running   0          63d      10.8.30.37
16 node37

```

南昌县环境部署问题解决

南昌县环境部署（深信服服务器），pod 没有网络。解决如下：

```

root@k8smaster01:/home/fs/k8sinstall# kubectl delete pods --selector=app=flannel -n=kube-system
pod "kube-flannel-ds-amd64-7n5rn" deleted
pod "kube-flannel-ds-amd64-8j2hv" deleted
pod "kube-flannel-ds-amd64-8qs4l" deleted
pod "kube-flannel-ds-amd64-8tbmm" deleted
pod "kube-flannel-ds-amd64-c2p89" deleted
pod "kube-flannel-ds-amd64-dlqjv" deleted
pod "kube-flannel-ds-amd64-hjnfq" deleted
pod "kube-flannel-ds-amd64-pnr97" deleted
pod "kube-flannel-ds-amd64-tlzmp" deleted
pod "kube-flannel-ds-amd64-vjrpk" deleted
pod "kube-flannel-ds-amd64-wwmsq" deleted
pod "kube-flannel-ds-amd64-xxwcx" deleted
root@k8smaster01:/home/fs/k8sinstall# kubectl get pods --selector=app=flannel -n=kube-system
NAME                READY     STATUS    RESTARTS   AGE
kube-flannel-ds-amd64-22j2r 1/1      Running   0          32s
kube-flannel-ds-amd64-47zbb 1/1      Running   0          24s
kube-flannel-ds-amd64-925ck 1/1      Running   0          23s
kube-flannel-ds-amd64-9bslc 1/1      Running   0          27s
kube-flannel-ds-amd64-fpv4h 1/1      Running   0          32s
kube-flannel-ds-amd64-g5v48 1/1      Running   0          23s
kube-flannel-ds-amd64-kj6zx 1/1      Running   0          24s
kube-flannel-ds-amd64-nd4ll 1/1      Running   0          29s
kube-flannel-ds-amd64-prkmb 1/1      Running   0          32s
kube-flannel-ds-amd64-qgdrh 1/1      Running   0          32s
kube-flannel-ds-amd64-wkc6x 1/1      Running   0          24s
kube-flannel-ds-amd64-x7xxg 1/1      Running   0          24s
root@k8smaster01:/home/fs/k8sinstall# kubectl -n lk exec emqx-deployment-6fdbf979b4-88zfr -it /bin/bash
kubectl exec [POD] [COMMAND] is DEPRECATED and will be removed in a future version. Use kubectl kubectl exec [POD] -- [COMMAND] instead.
bash-5.0$ sudo su
/opt/emqx # ping baidu.com
PING baidu.com (220.181.30.140): 56 data bytes
64 bytes from 220.181.30.140: seq=0 ttl=50 time=26.232 ms
64 bytes from 220.181.30.140: seq=1 ttl=50 time=26.214 ms
64 bytes from 220.181.30.140: seq=2 ttl=50 time=26.594 ms
^C
--- baidu.com ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 26.214/26.346/26.594 ms
/opt/emqx # ping 218.3.126.49
PING 218.3.126.49 (218.3.126.49): 56 data bytes
64 bytes from 218.3.126.49: seq=0 ttl=52 time=13.471 ms
64 bytes from 218.3.126.49: seq=1 ttl=52 time=12.156 ms
^C
--- 218.3.126.49 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 12.156/12.813/13.471 ms
/opt/emqx #

```