

Open Infrastructure - Task #6255

Find the right settings for kubernetes in ipv6 only settings

12/23/2018 12:57 PM - Nico Schottelius

Status:	Rejected	Start date:	12/23/2018
Priority:	Normal	Due date:	
Assignee:	Nico Schottelius	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
PM Check date:			

Description

Testing on

- 2a0a:e5c0:2:12:400:f0ff:fea9:c401
 - --pod-cidr 2a0a:e5c0:102:3::/64
 - --service-cidr 2a0a:e5c0:102:6::/64
- 2a0a:e5c0:2:12:400:f0ff:fea9:c402
 - --pod-cidr 2a0a:e5c0:102:4::/64
 - --service-cidr 2a0a:e5c0:102:7::/64
- 2a0a:e5c0:2:12:400:f0ff:fea9:c403
 - --pod-cidr 2a0a:e5c0:102:5::/64

Current findings below.

Follow up reading on:

- [calico/ipv6](#)
- [docker/ipv6 issue](#)
- [k8s/apt repo](#)

Possible options / next steps

- trying older docker version (< 17)
 - working around the docker/ipv6 issue
- using calico instead of bridge+host-local
 - working around the docker/ipv6 issue
- trying rkt instead of docker
 - working around the docker/ipv6 issue

401 / with most options

```
root@kube-master:~# kubeadm init --pod-network-cidr 2a0a:e5c0:102:3::/64 --apiserver-advertise-address=2a0a:e5c0:2:12:400:f0ff:fea9:c401 --service-cidr 2a0a:e5c0:102:6::/64
```

Result: api server not starting

402 / plain kubeadm init

- apiserver starts

403 / with --pod-network-cidr and bridge cni

```
root@kube-node2:~# kubeadm init --pod-network-cidr 2a0a:e5c0:102:5::/64
root@kube-node2:~# cat /etc/cni/net.d/10-bridge_v6.conf
```

```
{
  "cniVersion": "0.3.0",
  "name": "mynet",
  "type": "bridge",
  "bridge": "cbr0",
```

```

"isDefaultGateway": true,
"ipMasq": true,
"hairpinMode": true,
"ipam": {
  "type": "host-local",
  "ranges": [
    [
      {
        "subnet": "2a0a:e5c0:102:5::/64",
        "gateway": "2a0a:e5c0:102:5::1"
      }
    ]
  ]
}
}
}

```

- apiserver starts
- permission denied when trying to assign an IPv6 address
 - known bug in newer docker versions, which DISABLE ipv6 with a systctl!
- should be fixed in 0.7.x release of kubernetes-cni

k8s1 / service + pod cidr + calico

```

root@k8s1:~# kubeadm init --pod-network-cidr 2a0a:e5c0:102:3::/64 --service-cidr 2a0a:e5c0:102:6::/64

```

Result: failure at kubeadm init

k8s2 / podcidr + calico

Not yet changing the yaml files of calico, but finding out how far the setup goes w/o tuning/changing

```

kubeadm init --pod-network-cidr 2a0a:e5c0:102:4::/64

```

```

kubect1 apply -f https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/rbac-kdd.yaml

```

```

kubect1 apply -f https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/kubernetes-datastore/calico-networking/1.7/calico.yaml

```

Result:

- kubeadm init works
- Applying config for calico works
- calico pod is started
- Errors not finding nodename

k8s3 / podcidr + calico + calico guide

- following <https://docs.projectcalico.org/v3.4/getting-started/kubernetes/> instead of <https://kubernetes.io/docs/setup/independent/create-cluster-kubeadm/>
 - different versions, shows etcd for calico

After kubeadm init w/ pod cidr:

```

kubect1 apply -f \
https://docs.projectcalico.org/v3.4/getting-started/kubernetes/installation/hosted/etcd.yaml

```

```

kubect1 apply -f \
https://docs.projectcalico.org/v3.4/getting-started/kubernetes/installation/hosted/calico.yaml

```

Result: getting same error of missing nodename

k8s1+k8s2 / podcidr + calico + calico ipv6 guide

- based on <https://docs.projectcalico.org/v3.4/usage/ipv6>
- calico.yaml from <https://docs.projectcalico.org/v3.4/getting-started/kubernetes/installation/hosted/calico.yaml>

result:

- etcd need to have a service ip
- specifying --service-cidr makes kubeadm init fail

History

#1 - 12/23/2018 01:00 PM - Nico Schottelius

- Project changed from 45 to Open Infrastructure

- Subject changed from Find the right settings for kubernetes / ipv6 only to Find the right settings for kubernetes in ipv6 only settings

- Description updated

#2 - 12/23/2018 01:05 PM - Nico Schottelius

Next steps:

- Test with calico
- Test with kubernetes-cni >= 0.7.0

#3 - 12/23/2018 01:07 PM - Nico Schottelius

- Description updated

#4 - 12/23/2018 01:07 PM - Nico Schottelius

- Description updated

#5 - 12/23/2018 01:08 PM - Nico Schottelius

- Description updated

#6 - 12/23/2018 01:31 PM - Nico Schottelius

- Description updated

#7 - 12/23/2018 01:45 PM - Nico Schottelius

- Description updated

Using --service-cidr fails creating the cluster:

```
root@k8s1:~# kubeadm init --pod-network-cidr 2a0a:e5c0:102:3::/64 --service-cidr 2a0a:e5c0:102:6::/64
[init] Using Kubernetes version: v1.13.1
[preflight] Running pre-flight checks
  [WARNING Hostname]: hostname "k8s1" could not be reached
  [WARNING Hostname]: hostname "k8s1": lookup k8s1 on [2a0a:e5c0:2:1::5]:53: no such host
[preflight] Pulling images required for setting up a Kubernetes cluster
[preflight] This might take a minute or two, depending on the speed of your internet connection
[preflight] You can also perform this action in beforehand using 'kubeadm config images pull'
[kubelet-start] Writing kubelet environment file with flags to file "/var/lib/kubelet/kubeadm-flags.env"
[kubelet-start] Writing kubelet configuration to file "/var/lib/kubelet/config.yaml"
[kubelet-start] Activating the kubelet service
[certs] Using certificateDir folder "/etc/kubernetes/pki"
[certs] Generating "etcd/ca" certificate and key
[certs] Generating "etcd/healthcheck-client" certificate and key
[certs] Generating "etcd/server" certificate and key
[certs] etcd/server serving cert is signed for DNS names [k8s1 localhost] and IPs [2a0a:e5c0:2:12:400:f0ff:fea9:c401 127.0.0.1 ::1]
[certs] Generating "etcd/peer" certificate and key
[certs] etcd/peer serving cert is signed for DNS names [k8s1 localhost] and IPs [2a0a:e5c0:2:12:400:f0ff:fea9:c401 127.0.0.1 ::1]
[certs] Generating "apiserver-etcd-client" certificate and key
[certs] Generating "ca" certificate and key
[certs] Generating "apiserver-kubelet-client" certificate and key
[certs] Generating "apiserver" certificate and key
[certs] apiserver serving cert is signed for DNS names [k8s1 kubernetes kubernetes.default kubernetes.default.svc kubernetes.default.svc.cluster.local] and IPs [2a0a:e5c0:102:6::1 2a0a:e5c0:2:12:400:f0ff:fea9:c401]
[certs] Generating "front-proxy-ca" certificate and key
[certs] Generating "front-proxy-client" certificate and key
[certs] Generating "sa" key and public key
[kubeconfig] Using kubeconfig folder "/etc/kubernetes"
[kubeconfig] Writing "admin.conf" kubeconfig file
```

```
[kubeconfig] Writing "kubelet.conf" kubeconfig file
[kubeconfig] Writing "controller-manager.conf" kubeconfig file
[kubeconfig] Writing "scheduler.conf" kubeconfig file
[control-plane] Using manifest folder "/etc/kubernetes/manifests"
[control-plane] Creating static Pod manifest for "kube-apiserver"
[control-plane] Creating static Pod manifest for "kube-controller-manager"
[control-plane] Creating static Pod manifest for "kube-scheduler"
[etcd] Creating static Pod manifest for local etcd in "/etc/kubernetes/manifests"
[wait-control-plane] Waiting for the kubelet to boot up the control plane as static Pods from directory "/etc/kubernetes/manifests". This can take up to 4m0s
[kubelet-check] Initial timeout of 40s passed.
```

Unfortunately, an error has occurred:
timed out waiting for the condition

This error is likely caused by:

- The kubelet is not running
- The kubelet is unhealthy due to a misconfiguration of the node in some way (required cgroups disabled)

If you are on a systemd-powered system, you can try to troubleshoot the error with the following commands:

- 'systemctl status kubelet'
- 'journalctl -xeu kubelet'

Additionally, a control plane component may have crashed or exited when started by the container runtime. To troubleshoot, list all containers using your preferred container runtimes CLI, e.g. docker. Here is one example how you may list all Kubernetes containers running in docker:

```
- 'docker ps -a | grep kube | grep -v pause'
Once you have found the failing container, you can inspect its logs with:
- 'docker logs CONTAINERID'
error execution phase wait-control-plane: couldn't initialize a Kubernetes cluster
root@k8s1:~#
```

#8 - 12/23/2018 02:03 PM - Nico Schottelius

- Description updated

```
root@k8s2:~# kubectl apply -f https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/rbac-kdd.yaml
```

```
clusterrole.rbac.authorization.k8s.io/calico-node created
clusterrolebinding.rbac.authorization.k8s.io/calico-node created
root@k8s2:~#
```

```
root@k8s2:~# kubectl apply -f https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/kubernetes-datastore/calico-networking/1.7/calico.yaml
configmap/calico-config created
service/calico-typha created
deployment.apps/calico-typha created
poddisruptionbudget.policy/calico-typha created
daemonset.extensions/calico-node created
serviceaccount/calico-node created
customresourcedefinition.apiextensions.k8s.io/felixconfigurations.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/bgppeers.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/bgpconfigurations.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/ippools.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/hostendpoints.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/clusterinformations.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/globalnetworkpolicies.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/globalnetworksets.crd.projectcalico.org created
customresourcedefinition.apiextensions.k8s.io/networkpolicies.crd.projectcalico.org created
root@k8s2:~#
```

```
root@k8s2:~# kubectl get pods --all-namespaces
```

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
kube-system	calico-node-8lqc2	1/2	Running	0	44s
kube-system	coredns-86c58d9df4-7vrf7	0/1	ContainerCreating	0	6m49s
kube-system	coredns-86c58d9df4-gq54d	0/1	ContainerCreating	0	6m49s
kube-system	etcd-k8s2	1/1	Running	0	6m19s
kube-system	kube-apiserver-k8s2	1/1	Running	0	5m54s
kube-system	kube-controller-manager-k8s2	1/1	Running	0	5m52s
kube-system	kube-proxy-jr9qs	1/1	Running	0	6m49s
kube-system	kube-scheduler-k8s2	1/1	Running	0	6m5s

```
root@k8s2:~#
```

logs:

Dec 23 12:51:57 k8s2 kubelet[3670]: W1223 12:51:57.269589 3670 docker_sandbox.go:384] failed to read pod IP from plugin/docker: NetworkPlugin cni failed on the status hook for pod "coredns-86c58d9df4-gq54d_kube-system": CNI failed to retrieve network namespace path: cannot find network namespace for the terminated container "c681a4dc29f247dfd403900423514f855f8d676bbef3ce1ab5db16437379f2fc"

Dec 23 12:51:57 k8s2 kubelet[3670]: W1223 12:51:57.339750 3670 pod_container_deletor.go:75] Container "c681a4dc29f247dfd403900423514f855f8d676bbef3ce1ab5db16437379f2fc" not found in pod's containers

Dec 23 12:51:57 k8s2 kubelet[3670]: W1223 12:51:57.344057 3670 cni.go:302] CNI failed to retrieve network namespace path: cannot find network namespace for the terminated container "c681a4dc29f247dfd403900423514f855f8d676bbef3ce1ab5db16437379f2fc"

Dec 23 12:51:57 k8s2 dockerd[1361]: time="2018-12-23T12:51:57Z" level=info msg="shim reaped" id=b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb

Dec 23 12:51:57 k8s2 dockerd[1361]: time="2018-12-23T12:51:57.389778575Z" level=info msg="ignoring event" module=libcontainerd namespace=moby topic=/tasks/delete type="*events.TaskDelete"

Dec 23 12:51:57 k8s2 kubelet[3670]: E1223 12:51:57.528490 3670 remote_runtime.go:96] RunPodSandbox from runtime service failed: rpc error: code = Unknown desc = failed to set up sandbox container "b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb" network for pod "coredns-86c58d9df4-7vrf7_kube-system": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4-7vrf7_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/

Dec 23 12:51:57 k8s2 kubelet[3670]: E1223 12:51:57.529150 3670 kuberuntime_sandbox.go:65] CreatePodSandbox for pod "coredns-86c58d9df4-7vrf7_kube-system(851725b0-06b0-11e9-b3fa-0200f0a9c402)" failed: rpc error: code = Unknown desc = failed to set up sandbox container "b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb" network for pod "coredns-86c58d9df4-7vrf7_kube-system": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4-7vrf7_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/

Dec 23 12:51:57 k8s2 kubelet[3670]: E1223 12:51:57.529354 3670 kuberuntime_manager.go:662] createPodSandbox for pod "coredns-86c58d9df4-7vrf7_kube-system(851725b0-06b0-11e9-b3fa-0200f0a9c402)" failed: rpc error: code = Unknown desc = failed to set up sandbox container "b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb" network for pod "coredns-86c58d9df4-7vrf7_kube-system": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4-7vrf7_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/

Dec 23 12:51:57 k8s2 kubelet[3670]: E1223 12:51:57.529746 3670 pod_workers.go:190] Error syncing pod 851725b0-06b0-11e9-b3fa-0200f0a9c402 ("coredns-86c58d9df4-7vrf7_kube-system(851725b0-06b0-11e9-b3fa-0200f0a9c402)"), skipping: failed to "CreatePodSandbox" for "coredns-86c58d9df4-7vrf7_kube-system(851725b0-06b0-11e9-b3fa-0200f0a9c402)" with CreatePodSandboxError: "CreatePodSandbox for pod \"coredns-86c58d9df4-7vrf7_kube-system(851725b0-06b0-11e9-b3fa-0200f0a9c402)\" failed: rpc error: code = Unknown desc = failed to set up sandbox container \"b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb\" network for pod \"coredns-86c58d9df4-7vrf7_kube-system(851725b0-06b0-11e9-b3fa-0200f0a9c402)\" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/"

Dec 23 12:51:57 k8s2 dockerd[1361]: time="2018-12-23T12:51:57.661764208Z" level=info msg="No non-localhost DNS nameservers are left in resolv.conf. Using default external servers: [nameserver 8.8.8.8 nameserver 8.8.4.4]"

Dec 23 12:51:57 k8s2 dockerd[1361]: time="2018-12-23T12:51:57.670975343Z" level=info msg="Container 0f192031203f49f785746ef05fc572fc0032bcb0efa3f025d653a963852617d8 failed to exit within 2 seconds of signal 15 - using the force"

Dec 23 12:51:57 k8s2 dockerd[1361]: time="2018-12-23T12:51:57Z" level=info msg="shim docker-containerd-shim started" address="/containerd-shim/moby/470f83e9e58aede802f9401b35ba6b05769bb558528c27511e2e7ee4b15b60ba/shim.sock" debug=false pid=18198

Dec 23 12:51:57 k8s2 dockerd[1361]: time="2018-12-23T12:51:57Z" level=info msg="shim reaped" id=0f192031203f49f785746ef05fc572fc0032bcb0efa3f025d653a963852617d8

Dec 23 12:51:57 k8s2 dockerd[1361]: time="2018-12-23T12:51:57.844886065Z" level=info msg="ignoring event" module=libcontainerd namespace=moby topic=/tasks/delete type="*events.TaskDelete"

Dec 23 12:51:58 k8s2 dockerd[1361]: time="2018-12-23T12:51:58Z" level=info msg="shim docker-containerd-shim started" address="/containerd-shim/moby/57c70402308e302648a010281f872e8346f06f57cf47a6e593e1e7ae90697f81/shim.sock" debug=false pid=18290

Dec 23 12:51:58 k8s2 kubelet[3670]: E1223 12:51:58.356673 3670 cni.go:324] Error adding kube-system_coredns-86c58d9df4-gq54d/470f83e9e58aede802f9401b35ba6b05769bb558528c27511e2e7ee4b15b60ba to network calico/k8s-pod-network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/

Dec 23 12:51:58 k8s2 dockerd[1361]: time="2018-12-23T12:51:58Z" level=info msg="shim reaped" id=470f83e9e58aede802f9401b35ba6b05769bb558528c27511e2e7ee4b15b60ba

Dec 23 12:51:58 k8s2 dockerd[1361]: time="2018-12-23T12:51:58.653063452Z" level=info msg="ignoring event" module=libcontainerd namespace=moby topic=/tasks/delete type="*events.TaskDelete"

Dec 23 12:51:58 k8s2 kubelet[3670]: W1223 12:51:58.667299 3670 docker_sandbox.go:384] failed to read pod IP from plugin/docker: NetworkPlugin cni failed on the status hook for pod "coredns-86c58d9df4-7vrf7_kube-system": CNI failed to retrieve network namespace path: cannot find network namespace for the terminated container "b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb"

Dec 23 12:51:58 k8s2 kubelet[3670]: W1223 12:51:58.705493 3670 pod_container_deletor.go:75] Container "b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb" not found in pod's containers

Dec 23 12:51:58 k8s2 kubelet[3670]: W1223 12:51:58.729455 3670 cni.go:302] CNI failed to retrieve network namespace path: cannot find network namespace for the terminated container "b829a476b4875a4d8dd6973d470c6784c2c5544a5be80aef63e94b5ca56b51eb"

Dec 23 12:51:58 k8s2 kubelet[3670]: E1223 12:51:58.826461 3670 remote_runtime.go:96] RunPodSandbox from runtime service failed: rpc error: code = Unknown desc = failed to set up sandbox container "470f83e9e58aede802f9401b35ba6b05769bb558528c27511e2e7ee4b15b60ba" network for pod "coredns-86c58d9df4-gq54d": NetworkPlugin cni fa

```

iled to set up pod "coredns-86c58d9df4-gg54d_kube-system" network: stat /var/lib/calico/nodename: no such file
or directory: check that the calico/node container is running and has mounted /var/lib/calico/
Dec 23 12:51:58 k8s2 kubelet[3670]: E1223 12:51:58.827052      3670 kuberuntime_sandbox.go:65] CreatePodSandbox
for pod "coredns-86c58d9df4-gg54d_kube-system(8513a6d8-06b0-11e9-b3fa-0200f0a9c402)" failed: rpc error: code =
Unknown desc = failed to set up sandbox container "470f83e9e58aede802f9401b35ba6b05769bb558528c27511e2e7ee4b1
5b60ba" network for pod "coredns-86c58d9df4-gg54d": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4
-gg54d_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/n
ode container is running and has mounted /var/lib/calico/
Dec 23 12:51:58 k8s2 kubelet[3670]: E1223 12:51:58.827320      3670 kuberuntime_manager.go:662] createPodSandbox
for pod "coredns-86c58d9df4-gg54d_kube-system(8513a6d8-06b0-11e9-b3fa-0200f0a9c402)" failed: rpc error: code
= Unknown desc = failed to set up sandbox container "470f83e9e58aede802f9401b35ba6b05769bb558528c27511e2e7ee4b
15b60ba" network for pod "coredns-86c58d9df4-gg54d": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4
-gg54d_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/
node container is running and has mounted /var/lib/calico/

```

Checking container

```

root@k8s2:~# docker exec -ti k8s_calico-node_calico-node-8lqcz_kube-system_5e2f8e5e-06b1-11e9-b3fa-0200f0a9c40
2_3 ls /var/lib/calico/
root@k8s2:~#

```

-> indeed empty.

```

root@k8s2:~# wget https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/kubernet
e
s-datastore/calico-networking/1.7/calico.yaml
--2018-12-23 12:59:21-- https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/kub
er
netes-datastore/calico-networking/1.7/calico.yaml
Resolving docs.projectcalico.org (docs.projectcalico.org)... 2a03:b0c0:3:d0::d24:5001, 142.93.108.123
Connecting to docs.projectcalico.org (docs.projectcalico.org)|2a03:b0c0:3:d0::d24:5001|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 15075 (15K) [application/x-yaml]
Saving to: 'calico.yaml'

```

```

calico.yaml 100%[=====
] 14.72K --.-KB/s in 0.007s

```

2018-12-23 12:59:21 (2.12 MB/s) - 'calico.yaml' saved [15075/15075]

```

root@k8s2:~# wget https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/rbac-kdd.
yam
l
--2018-12-23 12:59:37-- https://docs.projectcalico.org/v3.3/getting-started/kubernetes/installation/hosted/rb
ac-
kdd.yaml
Resolving docs.projectcalico.org (docs.projectcalico.org)... 2a03:b0c0:3:d0::d24:5001, 142.93.108.123
Connecting to docs.projectcalico.org (docs.projectcalico.org)|2a03:b0c0:3:d0::d24:5001|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 1660 (1.6K) [application/x-yaml]
Saving to: 'rbac-kdd.yaml'

```

```

rbac-kdd.yaml 100%[=====
] 1.62K --.-KB/s in 0s

```

2018-12-23 12:59:37 (34.4 MB/s) - 'rbac-kdd.yaml' saved [1660/1660]

```

root@k8s2:~#

```

#9 - 12/23/2018 02:42 PM - Nico Schottelius

- Description updated

```

root@k8s3:~# kubectl apply -f \
> https://docs.projectcalico.org/v3.4/getting-started/kubernetes/installation/hosted/etcd.yaml
daemonset.extensions/calico-etcd created
service/calico-etcd created
root@k8s3:~# kubectl get pods --all-namespaces

```

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
kube-system	coredns-86c58d9df4-hhwvt	0/1	Pending	0	19m
kube-system	coredns-86c58d9df4-kbp2j	0/1	Pending	0	19m
kube-system	etcd-k8s3	1/1	Running	0	18m
kube-system	kube-apiserver-k8s3	1/1	Running	0	18m
kube-system	kube-controller-manager-k8s3	1/1	Running	0	18m
kube-system	kube-proxy-bzr7b	1/1	Running	0	19m
kube-system	kube-scheduler-k8s3	1/1	Running	0	18m

```

root@k8s3:~# kubectl apply -f \

```

```
> https://docs.projectcalico.org/v3.4/getting-started/kubernetes/installation/hosted/calico.yaml
configmap/calico-config created
secret/calico-etcd-secrets created
daemonset.extensions/calico-node created
serviceaccount/calico-node created
deployment.extensions/calico-kube-controllers created
serviceaccount/calico-kube-controllers created
clusterrole.rbac.authorization.k8s.io/calico-kube-controllers created
clusterrolebinding.rbac.authorization.k8s.io/calico-kube-controllers created
clusterrole.rbac.authorization.k8s.io/calico-node created
clusterrolebinding.rbac.authorization.k8s.io/calico-node created
root@k8s3:~#
```

```
oot@k8s3:~# kubectl get pods --all-namespaces
NAMESPACE      NAME                                                    READY   STATUS              RESTARTS   AGE
kube-system    calico-etcd-6pjx4                                       0/1     CrashLoopBackOff   1           29s
kube-system    calico-kube-controllers-5d94b577bb-rtrt9              0/1     Running           1           40s
kube-system    calico-node-mjq7j                                       0/1     CrashLoopBackOff   1           41s
kube-system    coredns-86c58d9df4-hhwvt                             0/1     ContainerCreating  0           20m
kube-system    coredns-86c58d9df4-kbp2j                             0/1     ContainerCreating  0           20m
kube-system    etcd-k8s3                                              1/1     Running           0           19m
kube-system    kube-apiserver-k8s3                                    1/1     Running           0           19m
kube-system    kube-controller-manager-k8s3                         1/1     Running           0           20m
kube-system    kube-proxy-bzr7b                                       1/1     Running           0           20m
kube-system    kube-scheduler-k8s3                                   1/1     Running           0           19m
root@k8s3:~#
```

logs:

```
Dec 23 13:38:45 k8s3 dockerd[1294]: time="2018-12-23T13:38:45Z" level=info msg="shim docker-containerd-shim started" address="/containerd-shim/moby/9660ec6261be39443fb39634cc57ee5213343cdef4740e9f3d4de7bce05315dd/shim.sock" debug=false pid=8790
Dec 23 13:38:45 k8s3 dockerd[1294]: time="2018-12-23T13:38:45Z" level=info msg="shim reaped" id=fcf7c79bf8cb66c677d41263ac13f267bbb9e49c05ef4e3cf665clf8bc696dff
Dec 23 13:38:45 k8s3 dockerd[1294]: time="2018-12-23T13:38:45.983006210Z" level=info msg="ignoring event" module=libcontainerd namespace=moby topic=/tasks/delete type="*events.TaskDelete"
Dec 23 13:38:46 k8s3 kubelet[15449]: E1223 13:38:46.157307 15449 remote_runtime.go:96] RunPodSandbox from runtime service failed: rpc error: code = Unknown desc = failed to set up sandbox container "fcf7c79bf8cb66c677d41263ac13f267bbb9e49c05ef4e3cf665clf8bc696dff" network for pod "coredns-86c58d9df4-hhwvt_kube-system": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4-hhwvt_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/
Dec 23 13:38:46 k8s3 kubelet[15449]: E1223 13:38:46.157394 15449 kuberuntime_sandbox.go:65] CreatePodSandbox for pod "coredns-86c58d9df4-hhwvt_kube-system(0794ed6d-06b5-11e9-8ab7-0200f0a9c403)" failed: rpc error: code = Unknown desc = failed to set up sandbox container "fcf7c79bf8cb66c677d41263ac13f267bbb9e49c05ef4e3cf665clf8bc696dff" network for pod "coredns-86c58d9df4-hhwvt_kube-system": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4-hhwvt_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/
Dec 23 13:38:46 k8s3 kubelet[15449]: E1223 13:38:46.157522 15449 kuberuntime_manager.go:662] createPodSandbox for pod "coredns-86c58d9df4-hhwvt_kube-system(0794ed6d-06b5-11e9-8ab7-0200f0a9c403)" failed: rpc error: code = Unknown desc = failed to set up sandbox container "fcf7c79bf8cb66c677d41263ac13f267bbb9e49c05ef4e3cf665clf8bc696dff" network for pod "coredns-86c58d9df4-hhwvt_kube-system": NetworkPlugin cni failed to set up pod "coredns-86c58d9df4-hhwvt_kube-system" network: stat /var/lib/calico/nodename: no such file or directory: check that the calico/node container is running and has mounted /var/lib/calico/
Dec 23 13:38:46 k8s3 kubelet[15449]: E1223 13:38:46.526783 15449 pod_workers.go:190] Error syncing pod cda5bd6d-06b7-11e9-8ab7-0200f0a9c403 ("calico-kube-controllers-5d94b577bb-rtrt9_kube-system(cda5bd6d-06b7-11e9-8ab7-0200f0a9c403)"), skipping: failed to "StartContainer" for "calico-kube-controllers" with CrashLoopBackOff: "Back-off 40s restarting failed container=calico-kube-controllers pod=calico-kube-controllers-5d94b577bb-rtrt9_kube-system(cda5bd6d-06b7-11e9-8ab7-0200f0a9c403)"
```

#10 - 12/23/2018 10:22 PM - Nico Schottelius

- Description updated

#11 - 12/23/2018 10:36 PM - Nico Schottelius

- Description updated

Modified IP in calico.yaml

```
root@k8s1:~/calico-34# kubectl apply -f calico.yaml
configmap/calico-config created
secret/calico-etcd-secrets created
daemonset.extensions/calico-node created
serviceaccount/calico-node created
deployment.extensions/calico-kube-controllers created
serviceaccount/calico-kube-controllers created
clusterrole.rbac.authorization.k8s.io/calico-kube-controllers created
clusterrolebinding.rbac.authorization.k8s.io/calico-kube-controllers created
clusterrole.rbac.authorization.k8s.io/calico-node created
clusterrolebinding.rbac.authorization.k8s.io/calico-node created
root@k8s1:~/calico-34#
```

```
root@k8s1:~/calico-34# ls
calico.yaml  etcd-v6.yaml  etcd.yaml
root@k8s1:~/calico-34# kubectl apply -f calico.yaml
configmap/calico-config created
secret/calico-etcd-secrets created
daemonset.extensions/calico-node created
serviceaccount/calico-node created
deployment.extensions/calico-kube-controllers created
serviceaccount/calico-kube-controllers created
clusterrole.rbac.authorization.k8s.io/calico-kube-controllers created
clusterrolebinding.rbac.authorization.k8s.io/calico-kube-controllers created
clusterrole.rbac.authorization.k8s.io/calico-node created
clusterrolebinding.rbac.authorization.k8s.io/calico-node created
root@k8s1:~/calico-34# kubectl get pods --all-namespaces
```

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
kube-system	calico-kube-controllers-5d94b577bb-jmzms	0/1	Running	0	36s
kube-system	calico-node-9d89x	0/1	Error	1	36s
kube-system	coredns-86c58d9df4-gh29q	0/1	ContainerCreating	0	5m30s
kube-system	coredns-86c58d9df4-h5fmp	0/1	ContainerCreating	0	5m30s
kube-system	etcd-k8s1	1/1	Running	0	4m53s
kube-system	kube-apiserver-k8s1	1/1	Running	0	4m59s
kube-system	kube-controller-manager-k8s1	1/1	Running	0	4m36s
kube-system	kube-proxy-jxp9f	1/1	Running	0	5m30s
kube-system	kube-scheduler-k8s1	1/1	Running	0	4m32s

```
root@k8s1:~/calico-34#
```

#12 - 12/23/2018 11:14 PM - Nico Schottelius

- Description updated

Only specifying the pod network without the service network stops us from changing the clusterip that is defined in calico.yaml:

```
root@k8s1:~/calico-34# kubectl apply -f etcd-v6.yaml
daemonset.extensions/calico-etcd created
The Service "calico-etcd" is invalid: spec.clusterIP: Invalid value: "2a0a:e5c0:102:3::edcd": provided IP is not in the valid range. The range of valid IPs is 10.96.0.0/12
```

However creating a cluster WITH specify the service-cidr fails (as seen above)

retrying with --service-cidr and looking for the exact error:

```
root@k8s1:~/calico-34# kubeadm init --pod-network-cidr 2a0a:e5c0:102:3::/64 --service-cidr 2a0a:e5c0:102:6::/64
4
```

logs:/debug

```
Dec 23 21:56:12 k8s1 kubelet[27221]: E1223 21:56:12.462137 27221 kubelet.go:2266] node "k8s1" not found
Dec 23 21:56:12 k8s1 kubelet[27221]: E1223 21:56:12.562383 27221 kubelet.go:2266] node "k8s1" not found
Dec 23 21:56:12 k8s1 kubelet[27221]: E1223 21:56:12.662695 27221 kubelet.go:2266] node "k8s1" not found
^C
```

```
root@k8s1:~# ping k8s1
PING k8s1 (k8s1 (2a0a:e5c0:2:12:400:f0ff:fea9:c401)) 56 data bytes
64 bytes from k8s1 (2a0a:e5c0:2:12:400:f0ff:fea9:c401): icmp_seq=1 ttl=64 time=0.081 ms
64 bytes from k8s1 (2a0a:e5c0:2:12:400:f0ff:fea9:c401): icmp_seq=2 ttl=64 time=0.063 ms
64 bytes from k8s1 (2a0a:e5c0:2:12:400:f0ff:fea9:c401): icmp_seq=3 ttl=64 time=0.071 ms
^C
```

```
--- k8s1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2044ms
rtt min/avg/max/mdev = 0.063/0.071/0.081/0.012 ms
root@k8s1:~#
```

```
root@k8s1:~# ps auxf | grep kubelet
root      27848  0.0  0.0 14856 1060 pts/2    S+   21:56   0:00          \_ grep --color=auto kubelet
root      27221  3.6  3.7 1336792 76476 ?        Ssl  21:56   0:01 /usr/bin/kubelet --bootstrap-kubeconfig=/etc/
kubernetes/bootstrap-kubelet.conf --kubeconfig=/etc/kubernetes/kubelet.conf --cluster-dns=2a0a:e5c0:102:6::a -
-config=/var/lib/kubelet/config.yaml --cgroup-driver=cgroupfs --network-plugin=cni --pod-infra-container-image
=k8s.gcr.io/pause:3.1 --resolv-conf=/run/systemd/resolve/resolv.conf
root@k8s1:~#
```

Full setup log/fail

```
root@k8s1:~/calico-34# kubeadm init --pod-network-cidr 2a0a:e5c0:102:3::/64 --service-cidr 2a0a:e5c0:102:6::/6
4
[init] Using Kubernetes version: v1.13.1
[preflight] Running pre-flight checks
[preflight] Pulling images required for setting up a Kubernetes cluster
[preflight] This might take a minute or two, depending on the speed of your internet connection
[preflight] You can also perform this action in beforehand using 'kubeadm config images pull'
[kubelet-start] Writing kubelet environment file with flags to file "/var/lib/kubelet/kubeadm-flags.env"
[kubelet-start] Writing kubelet configuration to file "/var/lib/kubelet/config.yaml"
[kubelet-start] Activating the kubelet service
[certs] Using certificateDir folder "/etc/kubernetes/pki"
[certs] Generating "etcd/ca" certificate and key
[certs] Generating "etcd/peer" certificate and key
[certs] etcd/peer serving cert is signed for DNS names [k8s1 localhost] and IPs [2a0a:e5c0:2:12:400:f0ff:fea9:
c401 127.0.0.1 ::1]
[certs] Generating "etcd/server" certificate and key
[certs] etcd/server serving cert is signed for DNS names [k8s1 localhost] and IPs [2a0a:e5c0:2:12:400:f0ff:fea
9:c401 127.0.0.1 ::1]
[certs] Generating "etcd/healthcheck-client" certificate and key
[certs] Generating "apiserver-etcd-client" certificate and key
[certs] Generating "ca" certificate and key
[certs] Generating "apiserver" certificate and key
[certs] apiserver serving cert is signed for DNS names [k8s1 kubernetes kubernetes.default kubernetes.default.
svc kubernetes.default.svc.cluster.local] and IPs [2a0a:e5c0:102:6::1 2a0a:e5c0:2:12:400:f0ff:fea9:c401]
[certs] Generating "apiserver-kubelet-client" certificate and key
[certs] Generating "front-proxy-ca" certificate and key
[certs] Generating "front-proxy-client" certificate and key
[certs] Generating "sa" key and public key
[kubeconfig] Using kubeconfig folder "/etc/kubernetes"
[kubeconfig] Writing "admin.conf" kubeconfig file
[kubeconfig] Writing "kubelet.conf" kubeconfig file
[kubeconfig] Writing "controller-manager.conf" kubeconfig file
[kubeconfig] Writing "scheduler.conf" kubeconfig file
[control-plane] Using manifest folder "/etc/kubernetes/manifests"
[control-plane] Creating static Pod manifest for "kube-apiserver"
[control-plane] Creating static Pod manifest for "kube-controller-manager"
[control-plane] Creating static Pod manifest for "kube-scheduler"
[etcd] Creating static Pod manifest for local etcd in "/etc/kubernetes/manifests"
[wait-control-plane] Waiting for the kubelet to boot up the control plane as static Pods from directory "/etc/
kubernetes/manifests". This can take up to 4m0s
[kubelet-check] Initial timeout of 40s passed.
```

Unfortunately, an error has occurred:

```
timed out waiting for the condition
```

This error is likely caused by:

- The kubelet is not running
- The kubelet is unhealthy due to a misconfiguration of the node in some way (required cgroups disable

d)

If you are on a systemd-powered system, you can try to troubleshoot the error with the following commands:

- 'systemctl status kubelet'
- 'journalctl -xeu kubelet'

Additionally, a control plane component may have crashed or exited when started by the container runtime.

To troubleshoot, list all containers using your preferred container runtimes CLI, e.g. docker.

Here is one example how you may list all Kubernetes containers running in docker:

- 'docker ps -a | grep kube | grep -v pause'

Once you have found the failing container, you can inspect its logs with:

- 'docker logs CONTAINERID'

```
error execution phase wait-control-plane: couldn't initialize a Kubernetes cluster
```

```
root@k8s1:~/calico-34#
```

Many such log messages on k8s1:

Dec 23 22:04:19 k8s1 kubelet[27221]: E1223 22:04:19.667280 27221 kubelet.go:2266] node "k8s1" not found

changing to k8s2 for re-testing

```
Dec 23 22:08:53 k8s2 kubelet[26960]: E1223 22:08:53.979215 26960 certificate_manager.go:348] Failed while requesting a signed certificate from the master: cannot create certificate signing request: Post https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:6443/apis/certificates.k8s.io/v1beta1/certificatesigningrequests: dial tcp [2a0a:e5c0:2:12:400:f0ff:fea9:c402]:6443: connect: connection refused
Dec 23 22:08:53 k8s2 kubelet[26960]: E1223 22:08:53.979265 26960 certificate_manager.go:269] Reached backoff limit, still unable to rotate certs: timed out waiting for the condition
Dec 23 22:08:53 k8s2 kubelet[26960]: E1223 22:08:53.998562 26960 kubelet.go:2266] node "k8s2" not found
```

```
root@k8s2:/etc/kubernetes# grep -ri 2a0a -r *
admin.conf:      server: https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:6443
controller-manager.conf:  server: https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:6443
kubelet.conf:      server: https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:6443
manifests/kube-controller-manager.yaml:  - --cluster-cidr=2a0a:e5c0:102:4::/64
manifests/etcd.yaml:     - --advertise-client-urls=https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:2379
manifests/etcd.yaml:     - --initial-advertise-peer-urls=https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:2380
manifests/etcd.yaml:     - --initial-cluster=k8s2=https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:2380
manifests/etcd.yaml:     - --listen-client-urls=https://127.0.0.1:2379,https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:2379
manifests/etcd.yaml:     - --listen-peer-urls=https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:2380
manifests/kube-apiserver.yaml:  - --advertise-address=2a0a:e5c0:2:12:400:f0ff:fea9:c402
manifests/kube-apiserver.yaml:  - --service-cluster-ip-range=2a0a:e5c0:102:7::/64
manifests/kube-apiserver.yaml:      host: 2a0a:e5c0:2:12:400:f0ff:fea9:c402
scheduler.conf:      server: https://[2a0a:e5c0:2:12:400:f0ff:fea9:c402]:6443
root@k8s2:/etc/kubernetes#
```

similar error. suspect probably resulting from non existing dns pods (?)

#13 - 01/02/2024 12:56 PM - Nico Schottelius

- Status changed from In Progress to Rejected