

Kubernetes, Containerd, and You

The fun I've encountered with using Kubernetes with ContainerD

- Initial Configuration

Initial Configuration

The Kubernetes documentation states:

=====

When using Docker, kubeadm will automatically detect the cgroup driver for the kubelet and set it in the `/var/lib/kubelet/config.yaml` file during runtime.

If you are using a different CRI, you must pass your `cgroupDriver` value to `kubeadm init`, like so:

```
apiVersion: kubelet.config.k8s.io/v1beta1
kind: KubeletConfiguration
cgroupDriver: <value>
```

Please mind, that you only have to do that if the cgroup driver of your CRI is not `cgroupfs`, because that is the default value in the kubelet already.

“ Note: Since `--cgroup-driver` flag has been deprecated by kubelet, if you have that in `/var/lib/kubelet/kubeadm-flags.env` or `/etc/default/kubelet` (`/etc/sysconfig/kubelet` for RPMs), please remove it and use the `KubeletConfiguration` instead (stored in `/var/lib/kubelet/config.yaml` by default).

Restarting the kubelet is required:

```
sudo systemctl daemon-reload
sudo systemctl restart kubelet
```

The automatic detection of cgroup driver for other container runtimes like CRI-O and containerd is work in progress.

- <https://kubernetes.io/docs/setup/production-environment/tools/kubeadm/install-kubeadm/#configure-cgroup-driver-used-by-kubelet-on-control-plane-node>

=====

Admitidly, I could not figure out how to pass that `cgroupDriver` value to `kubeadm`. Using the Arch Wiki, I did find that you can create / edit `/etc/kubernetes/kubelet.env` and add the line `KUBELET_EXTRA_ARGS="--cgroup-driver='systemd'"`. Then continue with the `kubeadm init`.

Once the init is done, the line can be removed from the `kubelet.env` file, and the line `cgroupDriver: systemd` can be added to the `/var/lib/kubelet/config.yaml` file.

- <https://wiki.archlinux.org/index.php/Kubernetes>

Note!

At the moment, you **WILL** need to add the line to the `/var/lib/kubelet/config.yaml` file of **EACH** node that you add to the cluster which is using a cgroup driver that is not cgroupfs. I know of the following:

- containerd
- CRI-O