

# How do I specify nvidia runtime from docker-compose.yml?

Asked 7 years, 3 months ago Modified 11 months ago Viewed 54k times



I am able to run a tensorflow container w/ access to the GPU from the command line w/ the following command

43

```
$ sudo docker run --runtime=nvidia --rm gcr.io/tensorflow/tensorflow:latest-gpu
```



I would like to be able to run this container from docker-compose. Is it possible to specify the `--runtime` flag from `docker-compose.yml`?



`docker` `docker-compose` `tensorflow` `nvidia-docker`

Share Improve this question Follow

asked Nov 24, 2017 at 2:51



rissem

553 1 5 11

## 4 Answers

Sorted by: Highest score (default)



Currently (Aug 2018), NVIDIA container runtime for Docker ([nvidia-docker2](#)) [supports Docker Compose](#).



By clicking "Accept all cookies", you agree Stack Exchange can store cookies on your device and disclose information in accordance with our [Cookie Policy](#).

Accept all cookies

Necessary cookies only

Customize settings

...dia to your GPU service. Docker

...e to show how to launch multiple GPU

[Share](#) [Improve this answer](#) [Follow](#)

answered Aug 28, 2018 at 18:04

[cedrickchee](#)

902 ● 8 ● 12

- 5 This option is possible only for docker-compose v2 for now. For more info, please look at the corresponding pull request here [github.com/docker/compose/pull/5405](https://github.com/docker/compose/pull/5405) – [hgasimov](#) Dec 30, 2018 at 9:33



This option is more elegant as it allows for better packaging using docker-compose – [Blaze](#) Jan 22, 2019 at 8:13

hm can different containers have different runtimes in one docker-compose file? – [kodlan](#) Feb 1, 2021 at 20:59

Do I still have to include the nvidia runtime in the daemon.json? – [Marcel](#) Sep 23, 2024 at 13:23



32

You should edit `/etc/docker/daemon.json`, adding the first level key `"default-runtime": "nvidia"`, restart docker daemon (ex. `"sudo service docker restart"`) and then all containers on that host will run with nvidia runtime.



More info on `daemon.json` [here](#)

[Share](#) [Improve this answer](#) [Follow](#)

answered Nov 26, 2017 at 11:58

[Gabriel Miretti aka](#)[gmiretti](#)

1,469 ● 13 ● 16



Or better: using `systemd` and assuming the path is `/usr/libexec/oci/hooks.d/nvidia`



By clicking "Accept all cookies", you agree Stack Exchange can store cookies on your device and disclose information in accordance with our [Cookie Policy](#).

```
-containers.conf <<EOF
```

```
runtime=nvidia
```

## Demo

Don't need to specify `--runtime=nvidia` since we set `default-runtime=nvidia` in the configuration step.

```
docker run --rm gcr.io/tensorflow/tensorflow:latest-gpu
```

Solution Inspired from my [tutorial about KATA runtime](#).

Share Improve this answer Follow

answered Jul 21, 2018 at 15:00



**Abdenmour TOUMI**

93.6k ● 42 ● 268 ● 269



Setting the default runtime seemed the most reasonable to me. However Gabriel's answer did not work for me.

1



Instead I found that the nvidia container toolkit can automatically configure the `daemon.json` [1].



```
sudo nvidia-ctk runtime configure --runtime=docker --set-as-default  
sudo service docker restart
```

This should be more future proof, since it now adds a different section to the `daemon.json`.

[1] <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/1.12.1/user-guide.html>

Share Improve this answer Follow


edited Mar 27, 2024 at 8:09

answered Mar 26, 2024 at 9:18




**CodeMonkey**

4,796 ● 2 ● 41 ● 51



By clicking "Accept all cookies", you agree Stack Exchange can store cookies on your device and disclose information in accordance with our [Cookie Policy](#).



By clicking "Accept all cookies", you agree Stack Exchange can store cookies on your device and disclose information in accordance with our [Cookie Policy](#).