
Subject: Re: CUDA support inside containers
Posted by [abufrejoval](#) on Fri, 18 Nov 2016 03:37:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

khorenko wrote on Wed, 16 November 2016 07:02May be `./NVIDIA-Linux-x86_64-XXX.run --no-kernel-module` option helps you?

<https://www.qubes-os.org/doc/install-nvidia-driver/>

First of all: Thanks for sticking with me! You guys give incredible support!

I'm guessing you are referring to this article: <http://scream.com/setting-cuda-linux-containers-2/> ..and the fact that there are pre-built docker containers actually available from Nvidia for CUDA.

I had been using the 'native' installers instead of the run file, so *.rpm for my CentoS7 container and *.deb for the Ubuntu one, which is why I didn't find that option right away.

If you use the run file, you can extract separate installers for the driver, the toolkit and the examples and there I can install along the lines of your recommendation and the article above.

But to make a long story short: LXC containers, even the so called unprivileged ones, get to see astonishing things from the host, including full dmesg info and--yes--/proc/modules, /proc/devices and I don't know what else...

That seems to make CUDA possible inside LXC (and Docker), while OpenVZ evidently hides much more, but either way, I have so far not been able to reproduce the results of the article above.

My guess is that the behaviour of the CUDA runtime has changed over time and that the 0.8 release may be behaving differently. My understanding is that it should really just use the device nodes to interact with the driver but from the system call traces I can see that it tries "to make things easy" for the user and will load the nvidia modules at application startup, if the user failed to do that himself (or automate that).

And all this logic which I can follow in the system call traces (attached) fails on OpenVZ, because the /proc file system misses information.

But it also fails on the plain CentOS7 machine with LXC containers using the same 0.8 software release, only it fails when it attempts to open the /dev/nvidia-uvm device node inside the container (outside works).

File permissions are alright but I guess LXC finally intervenes at this point.

I'll check with Docker on the CentOS control system next...

Having trouble attaching, will try in separate post!