Subject: Re: CUDA support inside containers
Posted by khorenko on Fri, 18 Nov 2016 18:07:02 GMT
View Forum Message <> Reply to Message

Ok, so what do we have at the moment (i did not try anything of this, so base on your words only and thus correct me if i miss something):
there are 2 ways of installation
1) via .run file which probably requires more steps but potentially allows to disable check for kernel module and thus allows to run CUDA program right now in OpenVZ Containers with no additional development

2) there is another way of CUDA software installation (btw, what's it? just rpm -ihv of some rpms? Which ones?)
It's easier, but it results in CUDA attempts to check for kernel module and fails dues to empty /proc/modules and inability to load module (the latter cannot be changed, otherwise anyone can load "bad" kernel module and crash the Hardware Node)

So i'm not against making things easier and don't see serious reasons why not to show /proc/modules content inside Container.

Do you have a great will to implement this and send the patch?

P.S. btw, try to workaround this check in order to understand how many of them are ahead: just save the output of /proc/modules on the host, copy the file inside the Container and mount the file onto /proc/modules file inside the Container.

```
# echo asdf > aaa
# cat aaa
asdf
# mount --bind aaa /proc/modules
# cat /proc/modules
asdf
```