
Subject: Connection to drivers located on the HN
Posted by [Balbuzard](#) on Wed, 29 Apr 2009 06:53:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,
I would like to provide the Graphic hardware into a VPS; I have seen that is may be dangerous (introduce security issues or damage the structure).
But I have seen that other solutions of virtualisation like VServer let as example the installation of Graphic Server inside a container.
So, first, which kind of issues can I encounter if I try to install a graphic server inside a container?
Why the installation of X-Server inside containers may be such harmful for the structure?

Then, OpenVZ lets its containers connect to the Internet, and so lets them to have a connection to the network hardware. I would like to use a driver (OpenGL) located on the HN (I am trying with solutions like X forwarding, or ssh -X but it does not work) Is there a way to establish a connection to the driver, the same way a connection is made to the network card to provide the Internet inside the VPE?

I know it may seem a little confused, so don't hesitate to ask for clarification if it is not clear enough;
Thanks for you help!

Subject: Re: Connection to drivers located on the HN
Posted by [maratrus](#) on Thu, 30 Apr 2009 09:44:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello Robert,

Quote:
Why the installation of X-Server inside containers may be such harmful for the structure?

The one of the main advantages provided by OpenVZ is an isolation, i.e. containers mustn't do anything harmful to HN. So, if the action that is done from inside the container and breaks the whole HN exists, it is a security hole which influences on the whole system's stability. So, we have to avoid all such actions and must eliminate the possibility of making our system inoperable. Giving a direct access to any hardware is always a risk. So, if we can avoid taking a risk, why don't we do that.
Moreover, talking about X-server, running X-server on the HN is not a limitation, all X-application are running inside VE not inside HN.

Quote:
I would like to use a driver (OpenGL) located on the HN (I am trying with solutions like X forwarding, or ssh -X but it does not work)

Could you please explain what OpenGL needs to run properly? May be we can workaround this issue? What error messages you got? Please, have a look at /var/log/messages or dmesg or use "strace" to find out the failed system call.

Anyway, don't hesitate filing a bug and describe what in particular you need and why do you need such configuration (it might be very important). Then you'll have a chance to get an explanation from more competent people than me.

Subject: Re: Connection to drivers located on the HN
Posted by [Balbuzard](#) on Fri, 01 May 2009 02:08:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ok, Thanks a lot for your answer,

OpenGL needs the direct rendering (DRI, which can be set up or down in the configuration file located at `/etc/X11/xorg.conf`), meaning that if an application needs OpenGL to work, it must be installed on the computer which has access to the direct rendering. So, I can not execute such applications in VPS because I have no way to connect directly to the Hardware and provide the direct rendering into the VPS.

By default, only one entity can have access to the direct rendering. In this case, the host has this access. So I have tried to disable it in the host and to enable it in the VPS, but it still does not work.

I will try to provide the access to the HW because it will provide the direct rendering in the container. I have seen the option `--devnodes` which can export the `/dev/devices` files. I just have to know which device running the graphic card I have to export to the container, and disable it in the HN. I will let you know what will happen

Thanks again for your answer!

Subject: Re: Connection to drivers located on the HN / graphics in container
Posted by [andreas2](#) on Fri, 04 Dec 2009 08:19:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hallo!

Did you find a working solution?

Is it possible to "move" the graphics device to the container,
if you only have one graphics card?

(Assuming direct output of HW-node and container messes it up?)

Is it possible to switch a device between hw-node and container (or different containers) while running (them)?

Perhaps it is possible to define a key (Alt+Fn) to allow automatic mapping (switching) of the device?

thanks, Andrew
