
Subject: enable the direct hardware access to the container

Posted by [Balbuzard](#) on Tue, 28 Apr 2009 00:46:15 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

I have read in the wiki-pages that the direct access to hardware is not available by default.

So is there a way to provides the direct access to the hardware nevertheless into the containers?

Thanks for your attention and your answers!

Subject: Re: enable the direct hardware access to the container

Posted by [vidhya](#) on Thu, 25 Jun 2009 17:50:40 GMT

[View Forum Message](#) <> [Reply to Message](#)

I'm curious to know how this is done too? How can the container be given direct access to the hardware? I'm assuming it read-only access. pls share your thoughts!

Subject: Re: enable the direct hardware access to the container

Posted by [maratrus](#) on Mon, 29 Jun 2009 06:59:54 GMT

[View Forum Message](#) <> [Reply to Message](#)

What kind of hardware you have in mind?

Subject: Re: enable the direct hardware access to the container

Posted by [Balbuzard](#) on Tue, 30 Jun 2009 05:39:49 GMT

[View Forum Message](#) <> [Reply to Message](#)

I have installed a heavy graphic application in the HN, and some non-graphic applications in VPS. But I would prefer the Graphic application to be installed in VPS instead of HN. In fact, I do not care if the graphic server or my graphic application crashes, as long as nothing else crashes with it. I need only the command line in the HN.

So, in a VPS, I have tried to install a graphic server, and a complete desktop, kde. This error has occurred during the installation: No userspace software suspend support in the kernel

The current kernel doesn't support userspace software suspend> Please recompile the kernel with the 'CONFIG_SOFTWARE_SUSPEND' option.

Once the installation finished, I have tried to stop the graphic server of the HN, then start it from inside this VE, but it does not recognize the hardware support.

I think I have to provide the direct access to the hardware, to the PCI-bus to reach the graphic card.

But I have read nothing to help me with that. I have read these discussions:

<http://forum.openvz.org/index.php?t=tree&th=7070&mid=34390&&rev=&reveal=> and

http://forum.openvz.org/index.php?t=rview&goto=35314#msg_35314

Thanks a lot for your help, any kind of trick/howto will be very welcomed!

Subject: Re: enable the direct hardware access to the container

Posted by [Balbuzard](#) on Wed, 01 Jul 2009 05:26:03 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hey,

I am still on this problem, here are some things I have tried:

I have tried to provide the VPS with the whole content of the /proc directory: `umount -l /vz/root/$VEID/proc`

`mount -t proc proc /vz/root/$VEID/proc`

and stopped the graphic server from the HN. Then, I entered this VPS and tried to start the graphic server by the `startx` command. The output returns: Fatal server error:

`xf86OpenConsole: Cannot open /dev/tty0 (no such file or directory) giving up.`

`xinit: Connection reset by peer (errno 104): unable to connect to X server`

`xinit: No such process (errno 3): Server error.`

So I have tried to provide the /dev/tty0 device with the command `vzctl set 301 --devnodes tty0:rw --save` but still the same error occurs, even after a restart of the VPS

I have no more ideas :-S

Thank you for your help!

Subject: Re: enable the direct hardware access to the container

Posted by [maratrus](#) on Wed, 01 Jul 2009 07:07:36 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello,

Quote:

I am still on this problem, here are some things I have tried:

please keep in mind that this is not an unstable but more that unstable workaround.

Quote:

So I have tried to provide the /dev/tty0 device with the command

Looks strange for me. Could you please remove tty* devices inside VE? And then invoke again

```
# vzctl set VEID --devnodes /dev/tty:rw
```

```
# vzctl set VEID --devnodes /dev/tty0:rw
```

Subject: Re: enable the direct hardware access to the container

Posted by [Balbuzard](#) on Thu, 02 Jul 2009 00:42:15 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi, thanks for your answer;

I have removed /dev/tty and /dev/tty0 in my VPS; and then added them again using the commands
vzctl set \$CTID --devnodes tty:rw --save
vzctl set \$CTID --devnodes tty0:rw --save
The shell returns setting devices saved parameters for CTID

Then, I have shut down the graphic server, and tried to start it up again from inside the VPS, the same error message is returned.

I have noticed that the rights of /dev/tty and /dev/tty0 are not the same in the HN and in the VPS, so I have set the rights in these devices of the VPS the same way they are in the HN, but still the same error message.

Does this configuration depend on the hardware on whom the openvz kernel is installed?

Thanks!

EDIT:

I have tried to move the devices tty0 and tty in the HN. Then, when I am trying to enter the VPS, vzctl returns;
Incorrect device name /dev/tty0: no such file or directory

Invalid value for DEVNODES=/tty0:rw, skipped
But the devices in VPS still appear! (And the same error again while trying to execute the command startx) Maybe it can help :-S

Subject: Re: enable the direct hardware access to the container

Posted by [Balbuzard](#) on Thu, 02 Jul 2009 04:45:48 GMT

[View Forum Message](#) <> [Reply to Message](#)

One other thing I don't understand, when you provided the whole content of /proc, is there any other daemon you have to install and run in the VPS?

I mean, the content of /proc is meant to provide communication to the Hardware, but is it enough to provide it to the VPS, or is there any link to provide to draw the liaison between the graphic server and the hardware?

Thanks for your answers!

Subject: Re: enable the direct hardware access to the container

Posted by [vidhya](#) on Wed, 08 Jul 2009 17:29:36 GMT

[View Forum Message](#) <> [Reply to Message](#)

I'm running linux 2.6.18 on the HN and FC-7.1 as a template for the VPS.

I am running a software in the container that basically gets the temperature of the disk, network and cpu from the /proc/diskstats folder from the HN. Now how do I allow the container access to get the real temperature of the HN?

Subject: Re: enable the direct hardware access to the container

Posted by [starfry](#) on Tue, 28 Jul 2009 14:59:19 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi Balbuzard, just seen your posting. I wanted to run my whole desktop in a VM so I could keep the HE as clean as possible.

I gave up.

I have since needed to run a heavy 32 bit 3D application on my 64 bit system. The way I achieved this was to use "schroot" to create a 32 bit container that is still part of the host, rather than a separate virtual machine. With proper configuration this accesses the hardware no problems at all.

Drop me a line if interested and I'll give you some pointers but it isn't an OpenVZ solution (sorry guys!).

I would still like to get a VPS with direct control over the graphics hardware to work, but I'm happy for now with how I have things.

Subject: Re: enable the direct hardware access to the container

Posted by [vidhya](#) on Wed, 29 Jul 2009 15:50:58 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi starfry,

Can you please explain how you made it work? I'm having a huge problem figuring out how i would populate the diskstats file in the proc directory in the VPS. I don't think mounting proc in the VPS is the right thing to do. Is there any way the container/VPS can discover its disk device, network ports etc, while it starts? If so, the diskstats file will automatically generated by the device driver.

Please share some pointers on how you enabled hardware access to the VPS.

Thanks,
Vidhya

Subject: Re: enable the direct hardware access to the container

Posted by [starfry](#) on Wed, 29 Jul 2009 17:23:47 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi Vidhya,

I don't know if you noticed in my post I said that I had it working but not using OpenVZ. I use a chroot built with "debootstrap" and I use "schroot" to access the chroot - it allows non root access and it configures bind mounts automatically. You do need 1.2 version of schroot, however, which I back-ported using "prevu".

I did this because I wanted to run the X-Plane flight sim which only comes as a 32 bit app. Initially I tried building a VPS but failed due to the need for direct access to the host's graphics driver, so I went the chroot route and it worked fine for my needs.

If you are interested I have documented the process on the X-Plane forum here:

<http://forums.x-plane.org/index.php?showtopic=34824&view=findpost&p=444528>

I guess this is all a little off topic for OpenVZ though