
Subject: Mysql Sockets?

Posted by [althalus](#) on Wed, 22 Apr 2009 06:26:08 GMT

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

Hi guys,

Experimenting with and testing OpenVZ at the moment, and so far I'm impressed. Things more or less just work, which is a much better experience with other virtualisation products I've tried.

I have a question regarding MySQL. I like the idea of having a mysql server that's more or less isolated from the network. Is it actually possible to share the socket between containers? I've seen a few approaches like this:

<http://forum.openvz.org/index.php?t=tree&goto=27708>

But it would seem that mounting a directory in the host does not actually mount the directory in the container.

I've seen other approaches that rely on a cronjob and a hard link, and whilst I have yet to test one of those, I'd prefer to not have the overhead of a cronjob running every 15 minutes and I'd rather not worry about databases being inaccessible in the time between the mysql socket changing and said cronjob being run again.

I realise this has probably been dealt with many a time, but there doesn't appear to be anything recent on the topic.

Subject: Re: Mysql Sockets?

Posted by [maratrus](#) on Wed, 22 Apr 2009 08:58:15 GMT

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

Hi,

Quote:

But it would seem that mounting a directory in the host does not actually mount the directory in the container.

no, that was the different and not OpenVZ-specific problem.

Files under /var/run directory may be removed at the beginning of the boot process. So, restarting a particular VE led to the removing files under /var/run/mysqld directory and in consequence of this action other VE were affected too, because /var/run/mysqld was a bind-mount.

Subject: Re: Mysql Sockets?

Posted by [althalus](#) on Wed, 22 Apr 2009 12:48:56 GMT

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

I meant the general approach, not the specific problem mentioned in that thread.

When I try to do:

```
mount -n --bind /var/lib/vz/shared/var_run_mysqld /var/lib/vz/root/101/var/run/mysqld
```

Even though the mount is present and working on the host, the files in the container do not reflect it at all. Using a mount point like this also seems to cause problems with vzquota. Perhaps I'm just missing a step somewhere?

Subject: Re: Mysql Sockets?

Posted by [maratrus](#) on Wed, 22 Apr 2009 14:22:30 GMT

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

Quote:

When I try to do:

```
mount -n --bind /var/lib/vz/shared/var_run_mysqld /var/lib/vz/root/101/var/run/mysqld
```

You have to do it after VE is mounted.

The best approach is using \$VEID.mount scripts.

Quote:

Even though the mount is present and working on the host, the files in the container do not reflect it at all.

I don't understand what you mean?

Quote:

Using a mount point like this also seems to cause problems with vzquota.

What problems you have in mind?

Subject: Re: Mysql Sockets?

Posted by [althalus](#) on Wed, 22 Apr 2009 23:14:42 GMT

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

Sorry, I wanted to include some more information in my previous reply, but did not have access to the machine.

Ok, so if mounting has to be done after the container starts, then logically unmounting must be done before the machine is unmounted? Ok, I wasn't sure on the order of those things, I haven't found anything that explains exactly what needs to be done in what steps. In that case, ignore my comment about vzquota. Using mount and unmount scripts makes sure everything happens in the right order?

```
mount -o bind /var/lib/vz/share /var/lib/vz/private/101/share
```

Works exactly as expected. The mount is exactly where I expect it. I can see the files of /var/lib/vz/share inside /var/lib/vz/private/101/share, as I expect to.

When I then do

```
vzctl enter 101
```

The /share folder is empty. If I create a file in this /share folder and then quit back out to the host machine, I can not see that file I just created in the container inside the share folder. I can stop and restart the container, and that file is still there, even though it's still not appearing in the /var/lib/vz/private/101/share folder on the host.

This seems contrary to how the file system seems to work with openvz...

BTW, running Ubuntu 8.04, and the latest versions of the openvz kernel and tools included in apt. (2.6.24-23-openvz for kernel)

EDIT: Ah, the files are appearing in /var/lib/vz/root... On closer examination I should have seen that the examples I was looking at were also putting the mount in root...

Ok, Can you explain the differences and purposes between the vz/private and vz/root directories? Both seem to be copies of the file system. If I wanted to edit something in a container from the host, in which cases would I make the changes where?

EDIT2: Can you define a default mount script that all machines will use automatically, or does there need to be a specific mount script (or link to a pre-existing one) for each machine?

EDIT3: What I want isn't possible short of applying a patch that seems like it could cause unwanted stability issues?

http://bugzilla.openvz.org/show_bug.cgi?id=985

Subject: Re: Mysql Sockets?

Posted by [maratrus](#) on Thu, 23 Apr 2009 06:15:33 GMT

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

Quote:

Using mount and unmount scripts makes sure everything happens in the right order?

Yes, a nice explanation is given in the "OpenVZ Users's Guide" (Managing OpenVZ scripts)

<http://download.openvz.org/doc/OpenVZ-Users-Guide.pdf>

Quote:

Ok, Can you explain the differences and purposes between the vz/private and vz/root directories?

Please read the following thread.

http://forum.openvz.org/index.php?t=msg&th=3988&#msg_20467

If you have any question, please ask I'll try to answer.

Quote:

What I want isn't possible short of applying a patch that seems like it could cause unwanted stability issues?

2.6.18 kernels should handle it properly I believe.

Subject: Re: Mysql Sockets?

Posted by [althalus](#) on Thu, 23 Apr 2009 23:14:11 GMT

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

maratrus wrote on Thu, 23 April 2009 02:15Quote:

Using mount and unmount scripts makes sure everything happens in the right order?

Yes, a nice explanation is given in the "OpenVZ Users's Guide" (Managing OpenVZ scripts)
<http://download.openvz.org/doc/OpenVZ-Users-Guide.pdf>

Ah, somehow I missed that.

Quote:

Quote:

Ok, Can you explain the differences and purposes between the vz/private and vz/root directories?

Please read the following thread.

http://forum.openvz.org/index.php?t=msg&th=3988&#msg_20467

If you have any question, please ask I'll try to answer.

Ok, that thread cleared up a lot of questions regarding the problems I've been running into. I understand how it all fits together a little better now. (So if I'm understanding correctly, changing anything in /vz/private is a Very Bad Idea as far as quota management, and so changing files with the VE turned off would subsequently be a Very Bad Idea.)

Quote:

Quote:

What I want isn't possible short of applying a patch that seems like it could cause unwanted stability issues?

2.6.18 kernels should handle it properly I believe.

Since that's obviously going backwards, is there anything I need to be aware of when downgrading the kernel?

EDIT: So I successfully grabbed the 2.6.18 kernel from the Git and compiled it.

The only bump in the road was that the make script didn't create the initrd.img for the 2.6.18 kernel. In the end, I unpacked the initrd image for the 2.6.24 kernel, replaced the directory in

lib/modules/ with the correct /lib/modules/ directory, and then repacked the new modified initrd.img.

So far everything seems to be working, and I have successfully shared the mysql socket of one machine with various others. No obvious problems, so far.

(As a side note, why does Ubuntu provide a DEV version of the OpenVZ kernel in their packages...? That seems to me a little.... odd, even if it is relatively stable... o.O)
