
Subject: fuse filesystem inside a container?

Posted by [George Georgalis](#) on Wed, 31 Aug 2011 17:22:11 GMT

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

Hi,

I'm relatively new to openvz but I think I got a good idea about managing containers, not looked much at the kernel side though. I'm trying to port an application running on a set of hosts, to containers running on the same hardware. The application requires both fuse and md devices/filesystems.

I'm able to make the md/fuse devices but I've been unable to configure a container to additionally include fuse filesystem capability.

```
vzctl set $id --devices b:9:2:rw c:10:229:rw --save
```

then after starting the containers,

```
vzctl exec $id mknod /dev/fuse c 10 229
```

```
vzctl exec $id mknod /dev/md2 b 9 2
```

I'm using the standard CentOS 5.6 template and running with the following rpms:

```
# rpm -aq | grep vz
```

```
vzquota-3.0.12-1
```

```
ovzkernel-devel-2.6.18-238.19.1.el5.028stab092.2
```

```
ovzkernel-2.6.18-238.19.1.el5.028stab092.2
```

```
vzctl-lib-3.0.28.3-1
```

```
vzctl-3.0.28.3-1
```

Fuse is being used to mount a gluster filesystem. Am I missing a step?

-George

--

George Georgalis, (415) 894-2710, <http://www.galis.org/>

Subject: fuse filesystem inside a container?

Posted by [George Georgalis](#) on Wed, 31 Aug 2011 17:38:10 GMT

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

Hi, (resending with the correct devices command)

I'm relatively new to openvz but I think I got a good idea about managing containers, not looked much at the kernel side though. I'm trying to port an application running on a set of hosts, to containers running on the same hardware. The application requires both fuse and

md devices/filesystems.

I'm able to make the md/fuse devices but I've been unable to configure a container to additionally include fuse filesystem capability.

```
vzctl set $id --devices b:9:2:rw --devices c:10:229:rw --save
```

then after starting the containers,

```
vzctl exec $id mknod /dev/fuse c 10 229  
vzctl exec $id mknod /dev/md2 b 9 2
```

I'm using the standard CentOS 5.6 template and running with the following rpms:

```
# rpm -aq | grep vz  
vzquota-3.0.12-1  
ovzkernel-devel-2.6.18-238.19.1.el5.028stab092.2  
ovzkernel-2.6.18-238.19.1.el5.028stab092.2  
vzctl-lib-3.0.28.3-1  
vzctl-3.0.28.3-1
```

Fuse is being used to mount a gluster filesystem. Am I missing a step?

-George

--

George Georgalis, (415) 894-2710, <http://www.galis.org/>

Subject: Re: fuse filesystem inside a container?

Posted by [Bogdan-Stefan Rotariu](#) on Wed, 31 Aug 2011 17:45:24 GMT

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

> Hi,
>

[...]

> I'm able to make the md/fuse devices but I've been unable to configure
> a container to additionally include fuse filesystem capability.

>
> vzctl set \$id --devices b:9:2:rw c:10:229:rw --save

>
> then after starting the containers,

>
> vzctl exec \$id mknod /dev/fuse c 10 229
> vzctl exec \$id mknod /dev/md2 b 9 2

>

[...]

>

> Fuse is being used to mount a gluster filesystem. Am I missing a step?

Yep,

Make sure you have the module in the container, kmod-fuse, and you can load it.

Subject: Re: fuse filesystem inside a container?

Posted by [George Georgalis](#) on Wed, 31 Aug 2011 18:37:05 GMT

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

On Wed, Aug 31, 2011 at 10:45 AM, Bogdan-Stefan Rotariu
<bogdan@rotariu.ro> wrote:

...

>> I'm able to make the md/fuse devices but I've been unable to configure
>> a container to additionally include fuse filesystem capability.

>>

>> vzctl set \$id --devices b:9:2:rw c:10:229:rw --save

>>

>> then after starting the containers,

>>

>> vzctl exec \$id mknod /dev/fuse c 10 229

>> vzctl exec \$id mknod /dev/md2 b 9 2

>>

>

> [...]

>

>>

>> Fuse is being used to mount a gluster filesystem. Am I missing a step?

>

>

> Yep,

>

> Make sure you have the module in the container, kmod-fuse, and you can
> load it.

Well that is certainly confusing, since the container doesn't have its own kernel. What does loading a kernel in a container mean?

The host has (and uses) the following modules:

/lib/modules/2.6.18-238.19.1.el5.028stab092.2/extra/fuse.ko

/usr/lib64/glusterfs/3.0.0/xlator/mount/fuse.so.0.0.0

I presume fuse.ko is the one loaded by the kernel while the gluster client uses fuse.so.0.0.0 in userspace.

In the container I have only,
/opt/glusterfs/3.2.3/lib64/glusterfs/3.2.3/xlator/mount/fuse .so.0.0.0

but when I try to load it I get:
insmod: error inserting
'/opt/glusterfs/3.2.3/lib64/glusterfs/3.2.3/xlator/mount/fuse.so.0.0.0':
-1 Operation not permitted

How do I load this fuse capability into the container? (nb the container gluster versions glusterfs-core-3.2.3-1 and glusterfs-fuse-3.2.3-1 are newer than the gluster running in the host, hopefully that won't be a issue through the migration to containers)

So, I'm trying to load the fuse and fuse-libs available in my rpm repos:

Downloading Packages:
fuse-2.7.4-8.el5.x86_64.rpm
Running rpm_check_debug
Running Transaction Test
Finished Transaction Test
Transaction Test Succeeded
Running Transaction
error: Couldn't fork %pre: Cannot allocate memory
error: install: %pre scriptlet failed (2), skipping fuse-2.7.4-8.el5

Installed:
fuse.x86_64 0:2.7.4-8.el5

Complete!

Downloading Packages:
fuse-libs-2.7.4-8.el5.x86_64.rpm
Running rpm_check_debug
Running Transaction Test
Finished Transaction Test
Transaction Test Succeeded
Running Transaction
Installing : fuse-libs
error: Couldn't fork %post: Cannot allocate memory

Installed:
fuse-libs.x86_64 0:2.7.4-8.el5

Complete!

With those "attempted" installs, I have:

```
# find /usr/lib* /lib* /opt -name \*fuse\* -type f
/lib64/libfuse.so.2.7.4
/opt/glusterfs/3.2.3/lib64/glusterfs/3.2.3/xlator/mount/fuse .so.0.0.0
```

Is there some kind of kernel module functionality I can load within a container? How do I go about that?

-George

--

George Georgalis, (415) 894-2710, <http://www.galis.org/>

Subject: Re: fuse filesystem inside a container?

Posted by [George Georgalis](#) on Wed, 31 Aug 2011 22:30:29 GMT

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

On Wed, Aug 31, 2011 at 11:37 AM, George Georgalis <george@galis.org> wrote:

> On Wed, Aug 31, 2011 at 10:45 AM, Bogdan-Stefan Rotariu

> <bogdan@rotariu.ro> wrote:

> ...

>>> I'm able to make the md/fuse devices but I've been unable to configure
>>> a container to additionally include fuse filesystem capability.

>>>

>>> vzctl set \$id --devices b:9:2:rw c:10:229:rw --save

>>>

>>> then after starting the containers,

>>>

>>> vzctl exec \$id mknod /dev/fuse c 10 229

>>> vzctl exec \$id mknod /dev/md2 b 9 2

>>>

>>

>> [...]

>>

>>>

>>> Fuse is being used to mount a gluster filesystem. Am I missing a step?

>>

>>

>> Yep,

>>

>> Make sure you have the module in the container, kmod-fuse, and you can
>> load it.

>

> Well that is certainly confusing, since the container doesn't have its
> own kernel. What does loading a kernel in a container mean?

>
> The host has (and uses) the following modules:
> /lib/modules/2.6.18-238.19.1.el5.028stab092.2/extra/fuse.ko
> /usr/lib64/glusterfs/3.0.0/xlator/mount/fuse.so.0.0.0
>
> I presume fuse.ko is the one loaded by the kernel while the gluster
> client uses fuse.so.0.0.0 in userspace.
>
> In the container I have only,
> /opt/glusterfs/3.2.3/lib64/glusterfs/3.2.3/xlator/mount/fuse .so.0.0.0
>
> but when I try to load it I get:
> insmod: error inserting
> '/opt/glusterfs/3.2.3/lib64/glusterfs/3.2.3/xlator/mount/fuse.so.0.0.0':
> -1 Operation not permitted
>
> How do I load this fuse capability into the container?

I've identified the warnings/errors in my prior email as memory allocation limits, so I set `privvmpages` to unlimited and reinstalled fuse & fuse-libs in the container without issue.

However I still have the confounding issue of how to I make the fuse filesystem available in the container `/proc/filesystems` like it is in the host?

-George

--

George Georgalis, (415) 894-2710, <http://www.galis.org/>
