Subject: Re: [patch 7/9] unprivileged mounts: allow unprivileged fuse mounts
Posted by Nigel Cunningham on Wed, 09 Jan 2008 09:29:24 GMT
View Forum Message <> Reply to Message

Hi.

Miklos Szeredi wrote:
>>>> On Tue 2008-01-08 12:35:09, Miklos Szeredi wrote:
>>>>> From: Miklos Szeredi <mszeredi@suse.cz>
>>>>>
>>>>> Use FS_SAFE for "fuse" fs type, but not for "fuseblk".
>>>>>
>>>>> FUSE was designed from the beginning to be safe for unprivileged users. This
>>>>> has also been verified in practice over many years. In addition unprivileged
>>>> Eh? So 'kill -9 no longer works' and 'suspend no longer works' is not
>>>> considered important enough to even mention?
>>> No. Because in practice they don't seem to matter. Also because
>>> there's no way in which fuse could be done differently to address
>>> these issues.
>> Could you clarify, please? I hope I'm getting the wrong end of the stick
>> - it sounds to me like you and Pavel are saying that this patch breaks
>> suspending to ram (and hibernating?) but you want to push it anyway
>> because you haven't been able to produce an instance, don't think
>> suspending or hibernating matter and couldn't fix fuse anyway?
>
> This patch has nothing to do with suspend or hibernate. What this
> patchset does, is help get rid of fusermount, a suid-root mount
> helper. It also opens up new possibilities, which are not fuse
> related.

That's what I thought. So what was Pavel talking about with "kill -9 no longer works" and "suspend no longer works" above? I couldn't understand it from the context.

> Fuse has bad interactions with the freezer, theoretically. In
> practice, I remember just one bug report (that sparked off this whole
> "do we need freezer, or don't we" flamefest), that actually got fixed
> fairly quickly, ...maybe. Rafael probably remembers better.

I think they just gave up and considered it unsolvable. I'm not sure it is.

>>> The 'kill -9' thing is basically due to VFS level locking not being
>>> interruptible. It could be changed, but I'm not sure it's worth it.
>>>
>>> For the suspend issue, there are also no easy solutions.
>> What are the non-easy solutions?
>
> The ability to freeze tasks in uninterruptible sleep, or more

> generally at any preempt point (except when drivers are poking
> hardware).

Couldn't some sort of scheduler based solution deal with the
uninterruptible sleeping case?

> I know this doesn't play well with userspace hibernate, and I don't
> think it can be resolved without going the kexec way.

I can see the desirability of kexec when it comes to avoiding the
freezer, but comes with its own problems too - having the original
context usable is handy, not having to set aside a large amount of space
for a second kernel is also desirable and there are still greater issues
of transferring information backwards and forwards between the two kernels.

Regards,

Nigel

_____

Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers