Subject: Re: [patch -mm 0/4] mqueue namespace
Posted by ebiederm on Fri, 20 Jun 2008 19:11:26 GMT
View Forum Message <> Reply to Message

"Serge E. Hallyn" <serue@us.ibm.com> writes:

>
> It is unfortunate that two actions are needed to properly complete the
> unshare, and we had definately talked about just using the mount before.
> I forget why we decided it wasn't practical, so maybe what you describe
> solves it...

What is worse, and I don't see a way around it: Is that we don't have
any callbacks to check where things are mounted. So we can't ensure the
proper kind of filesystem is mounted in the right place.

That is there is too much freedom in the mount apis to allow for reliable
operation.

> But at least the current patch reuses CLONE_NEWIPC for posix ipc, which
> also seems to make sense.

Sort of. I'm really annoyed with whoever did the posix mqueue support.
Adding the magic syscall that has to know the internal mount instead of
requiring the thing be mounted somewhere and just rejecting filedescriptors
for the wrong sorts of files.

Eric

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