Subject: Re: [RFC][PATCH 0/7] Clone PTS namespace
Posted by serge on Wed, 09 Apr 2008 19:16:34 GMT
View Forum Message <> Reply to Message

Quoting H. Peter Anvin (hpa@zytor.com):
> sukadev@us.ibm.com wrote:
>> We want to provide isolation between containers, meaning PTYs in container
>> C1 should not be accessible to processes in C2 (unless C2 is an ancestor).
>
> Yes, I certainly can understand the desire for isolation. That wasn't what
> my question was about.
>
>> The other reason for this in the longer term is for checkpoint/restart.
>> When restarting an application we want to make sure that the PTY indices
>> it was using is available and isolated.
>
> OK, this would be the motivation for index isolation.
>
>> A complete device-namespace could solve this, but IIUC, is being planned
>> in the longer term. We are hoping this would provide the isolation in the
>> near-term without being too intrusive or impeding the implementation of
>> the device namespace.
>
> I'm just worried about the accumulation of what feels like ad hoc
> namespaces, causing a very large combination matrix, a lot of which don't
> make sense.

Hmm, if we were to just call this CLONE_NEWDEV, would that (a) make
sense and (b) suitably address your (certainly valid) concern?

Basically for now CLONE_NEWDEV wouldn't yet be fully implemented, only
unsharing unix98 ptys...

-serge