

> Stanislav Kinsbursky<skinsbursky@parallels.com> writes:
>

>>> Stanislav Kinsbursky<skinsbursky@parallels.com> writes:
>>>

>>>> Stanislav Kinsbursky<skinsbursky@parallels.com> writes:
>>>>

>>>> Doing that independently of the rest of the sysctls is pretty horrible
>>>> and confusing to users. What I am planning might suit your needs and
>>>> if not we need to talk some more about how to get the vfs to do
>>>> something reasonable.
>>>>

>>>>
>>>> Ok, Eric. Would be glad to discuss your sysctls plans.
>>>> But actually you already know my needs: I would like to make sysctls work in the
>>>> way like sysfs does: i.e. content of files depends on mount maker -
>>>> not viewer.

>>>
>>> What drives the desire to have sysctls depend on the mount maker?
>>

>> Because we can (will, actually) have nested fs root's for containers. IOW,
>> container's root will be accessible from it's creator context. And I want to
>> tune container's fs from creators context.

>
> Tuning the child context from the parent context is an entirely
> reasonable thing to do. To affect a namespace that is not yours
> the requirement is simply that we don't use current to lookup the
> sysctl. So what I am proposing should work for your case.
>

Could you explain, what are you proposing?
I still don't know any details about it.

>>> Especially what drives that desire not to have it have a /proc/<pid>/sys
>>> directory that reflects the sysctls for a given process.
>>>

>>
>> This is not so important for me, where to access sysctl's. But I'm worrying
>> about backward compatibility. IOW, I'm afraid of changing path
>> "/proc/sys/sunrpc/" to "/proc/<pid>/sys/sunrpc". This would break a lot of
>> user-space programs.
>

> The part that keeps it all working is by adding a symlink from /proc/sys
> to /proc/self/sys. That technique has worked well for /proc/net, and I
> don't expect there will be any problems with /proc/sys either. It is
> possible but is very rare for the introduction of a symlink in a path
> to cause problems.
>

Probably I don't understand you, but as I see it now, symlink to "/proc/self/"
is unacceptable because of the following:

- 1) will be used current context (any) instead of desired one
- 1) if CT has other pid namespace - then we just have broken link.

> Eric

>

--

Best regards,
Stanislav Kinsbursky
