Subject: Re: [PATCH 01/11] SYSCTL: export root and set handling routines
Posted by Stanislav Kinsbursky on Tue, 10 Jan 2012 10:38:07 GMT
View Forum Message <> Reply to Message

> 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.

> 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/sunprc/*" to "/proc/<pid>/sys/sunrpc". This would break a lot of
user-space programs.

--
Best regards,
Stanislav Kinsbursky