
Subject: Re: [PATCH] namespaces: uts_ns: make information visible
via /proc/PID/uts directory

Posted by [Sam Vilain](#) on Mon, 22 May 2006 23:07:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

Alan Cox wrote:

>On Llu, 2006-05-22 at 17:24 +1200, Sam Vilain wrote:

>

>

>>From: Sam Vilain <sam.vilain@catalyst.net.nz>

>>

>>Export the UTS information to a per-process directory /proc/PID/uts,
>>that has individual nodes for hostname, ostype, etc - similar to
>>those in /proc/sys/kernel

>>

>>

>

>Can you explain the locking being used here against the name being
>changed at the same moment ?

>

Is this a test? :-)

Let's see, get_task_pid locks the task struct (so that it doesn't go
away while we're de-referencing the nsproxy and uts_ns etc), and the
kobj references are assumed to be enough to avoid the references
dropping away in the meantime.

I didn't grab uts_sem. That semaphore could be made per-uts_ns, in
theory. Whether anyone cares about contention that much is another
question.

I intended this to be exploratory, it wasn't really mergable. Should
have added an [RFC] tag, sorry about that.

Sam.
