

---

Subject: Re: [RFC][PATCH 4/5] utsname namespaces: sysctl hack

Posted by [ebiederm](#) on Wed, 19 Apr 2006 15:52:42 GMT

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

---

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

> Quoting Kirill Korotaev (dev@sw.ru):

>> Serge,

>>

>> can we do nothing with sysctls at this moment, instead of committing hacks?

>

> Please look closer at the patch.

>

> I \*am\* doing nothing with sysctls.

>

> system\_utsname no longer exists, and the way to get to that is by using

> init\_uts\_ns.name. That's all this does.

Ack. I probably read that question backwards.

Yes you must at least touch kernel/sysctl.c when you kill  
system\_utsname.

I read it as: Can we do nothing better with sysctls than committing hacks?

Eric

---