
Subject: Re: [RFC][PATCH 1/2] Virtualization of UTS
Posted by [ebiederm](#) on Fri, 24 Mar 2006 19:50:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev <dev@sw.ru> writes:

>>> This patch introduces utsname namespace in system, which allows to have
>>> different utsnames on the host.
>>> Introduces config option CONFIG_UTS_NS and uts_namespace structure for this.
>> Ok. It looks like we need to resolve the sysctl issues before we merge
>> either patch, into the stable kernel.
> I disagree with you. Right now we can have sysctl and proc for init namespaces
> only.
> And when sysctl and proc are virtualized somehow, we can fix all these.
> I simply don't expect /proc and sysctl to be done quickly. As we have very
> different approaches. And there is no any consensus. Why not to commit
> working/agreed parts then?

So getting this code into Andrews development tree (as long as he is willing to accept it) looks very reasonable. We can't change the interface once we get into the stable kernel because that becomes part of the ABI.

So all I am saying is that this code is clearly not yet ready for the stable branch, because we plan to change the sysctl interface.

>> We also need to discuss the system call interface, as without one
>> the functionality is unusable :)
> I also don't see why it can be separated. There is an API in namespaces, and how
> it is mapped into syscalls is another question. At least it doesn't prevent us
> from committing virtualization itself, agree?

Separating the patches makes a lot of sense. Putting something into the kernel without any in tree users is a problem.

Eric
