
Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by [ebiederm](#) on Wed, 07 Mar 2007 23:13:15 GMT

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

Dave Hansen <hansendc@us.ibm.com> writes:

> On Wed, 2007-03-07 at 15:59 -0600, Serge E. Hallyn wrote:
>> Space saving was the only reason for nsproxy to exist.
>>
>> Now of course it also provides the teensiest reduction in # instructions
>> since every clone results in just one reference count inc for the
>> nsproxy rather than one for each namespace.
>
> If we have 7 or 8 namespaces, then it can save us a significant number
> of atomic instructions on each of the refcounts, plus touching all of
> the cachelines, etc...

Well we still have a global lock on the fork path so there is only so much we can do to improve things. The global process list, and there are some interesting posix signal handling rules that limit how much we can relax that restriction.

However with namespaces we have a natural limit on how many we will have. There aren't that many spaces for global names.

I don't know the situation well enough for resource controllers but I suspect we might not have any kind of natural limit (except what a single person can comprehend) to the kind of resource we will ultimately want to control which tends to imply we will have more of those.

Eric

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
