Eric W. Biederman wrote:

>>> as I said, I'd opt for having a new clone() syscall in
>>> addition to the existing one, with a separate 64bit
>>> set of flags to decide what namespaces should be created
>>> or cloned. there is no problem with putting 'important'
>>> or generally 'useful' flags (like for example for pid,
>>> uts or lightweight network isolation) into the existing
>>> clone call (will require a simple mapping if done properly)
>>> so that they can be used with 'older' libc interfaces too
>>>
>>> I know, it would be 'nice' to keep the existing clone()
>>> interface, but I think it already has become a complication
>>> we should avoid (and we have not even used up all the
>>> available flags :)
>> agree and so does Kirill.
>>
>>> are there any strong arguments against having a new
>>> clone() syscall, which I was missing so far?
>> I don't see any.
>>
>> I'm going to revive execns() syscall into a clone_ns() syscall as suggested
>> by Kirill and you. Then, others will be free to nack ;)
>
> I think it is silly, but I see not real problems with the idea.

that's not a violent agreement :)

i'll work on it.

thanks,

C.