
Subject: Re: [RFC][PATCH 1/2] add user namespace [try #2]
Posted by [Cedric Le Goater](#) on Mon, 11 Sep 2006 08:46:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev wrote:

```
>> On Thu, Sep 07, 2006 at 08:05:30PM +0400, Kirill Korotaev wrote:
>>
>>> BTW...
>>>
>>>
>>>> --- 2.6.18-rc4-mm3.orig/include/linux/sched.h
>>>> +++ 2.6.18-rc4-mm3/include/linux/sched.h
>>>> @@ -26,6 +26,7 @@
>>>> #define CLONE_STOPPED 0x02000000 /* Start in stopped state */
>>>> #define CLONE_NEWUTS 0x04000000 /* New utsname group? */
>>>> #define CLONE_NEWIPC 0x08000000 /* New ipc */
>>>> + #define CLONE_NEWUSER 0x10000000 /* New user */
>>> we have place for 3 namespaces more only.
>>> Does anyone have a plan what to do then?
>>
>> what about having a new clone syscall with 32 or
>> better 64 bits reserved for namespace stuff, and
>> only put basic/generic namespaces or even aggregate
>> flags into the existing clone interface?
>>
>> something like: uts+ipc+user -> CLONE_NEWXYZ
>> but CLONE2_NEWUTS, CLONE2_NEWIPC, CLONE2_NEWUSER
> I would suggest to do it another way then:
> remove CLONES_NEWXXXNS from clone() at all (except for MNT NS for compatibility)
> and introduce sys_clone_ns() with totatally new 64bit flags like
> CLONE_NS_UTS
> CLONE_NS_IPC
> CLONE_NS_USER
> CLONE_NS_NET
```

yep. I like the idea of a specific syscall. It would certainly help us to handle some corner cases in the namespaces.

OTOH, the unshare/clone semantic is right in most cases.

How would the community feel about this ? would they say "fix unshare/clone" or this is a new API, move it somewhere else ?

thanks,

C.

Containers mailing list

