
Subject: Re: [PATCH 1/2] Extend sys_clone and sys_unshare system calls API
Posted by [Cedric Le Goater](#) on Thu, 24 Jan 2008 17:09:56 GMT

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

Pavel Machek wrote:

> On Wed 2008-01-16 15:58:55, Pavel Emelyanov wrote:

>> There's only one bit in the clone_flags left, so we won't be able
>> to create more namespaces after we make it busy. Besides, for
>> checkpoint/restart jobs we might want to create tasks with given
>> pids (virtual of course). And nobody knows for sure what else might
>> be required from clone() in the future.

>>
>> This is an attempt to create a extendable API for clone and unshare.
>> Actually this patch is a request for comment about the overall
>> design. If it will turn out to "look good", then we'll select some
>> better names for new flag and data types.

>>
>> I use the last bit in the clone_flags for CLONE_LONGARG. When set it
>> will denote that the child_tidptr is not a pointer to a tid storage,
>> but the pointer to the struct long_clone_struct which currently
>> looks like this:

```
>> struct long_clone_arg {  
>> int size;  
>> };
```

>
> Ugly as night, I'd say. (AI said it better). What about just adding
> clone2 syscall, that takes u64?

yes but we would need more something like :

```
long sys_clone64(unsigned long flags_high, unsigned long flag_low)
```

if we want the syscall to be supported on 32bit arch. clone2 is also
being used on ia64 already.

C.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
