

---

Subject: Re: [PATCH 1/2] Extend sys\_clone and sys\_unshare system calls API  
Posted by [Pavel Machek](#) on Wed, 23 Jan 2008 20:59:30 GMT

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

---

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. (Al said it better). What about just adding  
clone2 syscall, that takes u64?

--

(english) <http://www.livejournal.com/~pavelmachek>

(cesky, pictures) <http://atrey.karlin.mff.cuni.cz/~pavel/picture/horses/blog.html>

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---