

---

Subject: Re: [PATCH] An attempt to have an unlimitedly extendable sys\_clone  
Posted by [Daniel Hokka Zakrisso](#) on Wed, 16 Jan 2008 08:07:21 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I wrote:

> Pavel Emelyanov wrote:

>> [snip]

>>

>>>> +static struct long\_clone\_arg \*get\_long\_clone\_arg(int \_\_user

>>>> \*child\_tidptr)

>>>> +{

>>>> + int size;

>>>> + struct long\_clone\_arg \*carg;

>>>> +

>>>> + if (get\_user(size, child\_tidptr))

>>>> + return ERR\_PTR(-EFAULT);

>>>> +

>>>> + if (size > sizeof(struct long\_clone\_arg))

>>>> + return ERR\_PTR(-EINVAL);

>>>

>>> This means that software built against a newer kernel won't work on an

>>> older one. Surely that's not intended?

>>

>> It is intended. If I ask an old kernel to clone the mq namespace, but

>> it doesn't support such, that I'd better like to get an -EINVAL error

>> rather than be silently held in an old global namespace.

>

> That rules out using the struct for things like child\_tidptr, the desired

> pid for C/R, etc. I think the version Dave Hansen proposed would be

> better, or if it's really just for bits, use an array rather than a struct

> to make that obvious.

s/Dave Hansen/Oren Laadan/

Sorry, too early for me.

>> [snip]

>>

--

Daniel Hokka Zakrisson

---

Containers mailing list

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

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

---