

---

Subject: Re: [PATCH 0/3] clone64() and unshare64() system calls  
Posted by [Cedric Le Goater](#) on Thu, 10 Apr 2008 06:48:50 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

H. Peter Anvin wrote:

> sukaddev@us.ibm.com wrote:

>> This is a resend of the patch set Cedric had sent earlier. I ported  
>> the patch set to 2.6.25-rc8-mm1 and tested on x86 and x86\_64.

>> ---

>>

>> We have run out of the 32 bits in clone\_flags !

>>

>> This patchset introduces 2 new system calls which support 64bit  
>> clone-flags.

>>

>> long sys\_clone64(unsigned long flags\_high, unsigned long flags\_low,  
>> unsigned long newsp);

>>

>> long sys\_unshare64(unsigned long flags\_high, unsigned long  
>> flags\_low);

>>

>> The current version of clone64() does not support CLONE\_PARENT\_SETTID  
>> and CLONE\_CHILD\_CLEARPID because we would exceed the 6 registers limit  
>> of some arches. It's possible to get around this limitation but we  
>> might not  
>> need it as we already have clone()

>>

>

> I really dislike this interface.

>

> If you're going to make it a 64-bit pass it in as a 64-bit number,  
> instead of breaking it into two numbers. Better yet, IMO, would be to  
> pass a pointer to a structure like:

>

> struct shared {  
> unsigned long nwords;  
> unsigned long flags[];  
> };

>

> ... which can be expanded indefinitely.

ok.

What about the copy\_from\_user() overhead ? is this something we care  
about for a clone like syscall ?

If not, this would certainly make our life simpler to extend clone flags.  
I'm ready to implement anything if someone would just tell me in which

direction.

Thanks !

C.

---

Containers mailing list

Containers@lists.linux-foundation.org

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

---