
Subject: Re: [PATCH] An attempt to have an unlimitedly extendable sys_clone
Posted by [Dave Hansen](#) on Wed, 16 Jan 2008 17:52:57 GMT

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

On Wed, 2008-01-16 at 10:26 +0300, Pavel Emelyanov wrote:

> > 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.

There's at least one thing that this doesn't handle. Let's say we have
a clone64(long, long) but that the kernel we're running is old and don't
know about any valid flags in the second long, yet.

Somebody who knows about lots of new flags 10 years down the line passes
clone(0x1234, 0x0) to the old kernel. That's OK, because the kernel
knows about all of the flags that *ARE* set.

Your proposal is basically this:

```
clone(flags, {size, stuff});
```

and if the in-kernel size isn't exactly the same as the size passed in,
it punts. It doesn't matter if anything in "stuff" was actually going
to be used, or if the old kernel *could* have handled it, it just punts.

Honestly, I don't think it's that big of a deal, I just wish you'd
acknowledge it. ;)

-- Dave

Containers mailing list

Containers@lists.linux-foundation.org

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