
Subject: Re: [PATCH 0/3] clone64() and unshare64() system calls
Posted by [hpa](#) on Thu, 10 Apr 2008 18:31:41 GMT

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

sukadev@us.ibm.com wrote:

> |
> | I thought that the consensus was that adding a new system call was
> | better than trying to force extensibility on to the existing
> | non-extensible system call.
>
> There were couple of objections to extensible system calls like
> sys_indirect() and to Pavel's approach.
>

This is a very different thing, though. sys_indirect is pretty much a mechanism for having a sideband channel -- a second ABI -- into each and every system call, making it extremely hard to analyze what the full set of impact of a specific system call is. Worse, as it was being proposed to have been used, it would have set state variables inside the kernel in a very opaque manner.

> | But if we are adding a new system call, why not make the new one
> | extensible to reduce the need for yet another new call in the future?
>
> hypothetically, can we make a variant of clone() extensible to the point
> of requiring a copy_from_user() ?

The only issue is whether or not it's acceptable from a performance standpoint. clone() is reasonably expensive, though.

-hpa

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