
Subject: Extending syscalls (was: [PATCH 1/2] Extend sys_clone and sys_unshare system calls API)

Posted by [corbet](#) on Thu, 17 Jan 2008 15:02:05 GMT

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

Al Viro sez:

> Nah, just put an XML parser into the kernel to have the form match the
> contents...

>

> Al "perhaps we should newgroup alt.tasteless.api for all that stuff" Viro

Heh, indeed. But we do seem to have a recurring problem of people wanting to extend sys_foo() beyond the confines of its original API.

I've observed a few ways of doing that:

- create sys_foo2() (or sys_foo64(), or sys_fooat(), or sys_pfoo(), or...) and add the new stuff there.
- Put a version number into the API somewhere - wireless extensions, for example.
- Set a flag saying "I've stashed some additional parameters somewhere else." That's sys_indirect() and the current proposal for extending clone().
- Just do it all with a kernel-based XML parser. I think we should call this approach sys_viro() in honor of its champion.
- Do it all in sysfs

The first approach has traditionally been the most popular. If we have a consensus that this is the way to extend system calls in the future, it would be nice to set that down somewhere. We could avoid a lot of API blind alleys that way.

jon

Containers mailing list

Containers@lists.linux-foundation.org

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