

## Re: how do you call userspace syscalls (e.g. sys\_rename) from inside kernel

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-10/2294.html>

---

**From:** Luke Kenneth Casson Leighton (*lkcl\_at\_lkcl.net*)

**Date:** 10/08/04

Date: Fri, 8 Oct 2004 16:35:23 +0100  
To: Fabiano Ramos <ramos\_fabiano@yahoo.com.br>

On Fri, Oct 08, 2004 at 10:07:04AM -0300, Fabiano Ramos wrote:

> *On Fri, 2004-10-08 at 14:04 +0100, Luke Kenneth Casson Leighton wrote:*  
> > *could someone kindly advise me on the location of some example code in*  
> > *the kernel which calls one of the userspace system calls from inside the*  
> > *kernel?*  
> >  
> > *alternatively if this has never been considered before, please could*  
> > *someone advise me as to how it might be achieved?*  
> >  
> >  
> > *you cannot do that. For every sys\_xx there is a do\_xx, that can*  
> > *be called from inside the kernel.*

so, there's a do\_rename (yes i found that and ISTRC that when i used it i can't exactly remember what the problem was: either i got an error code -14 or i got "warning symbol do\_rename not found" when my module was linked together, even though it says EXPORT\_SYMBOL(do\_rename) in fs/namei.c, so i was forced to cut/paste sys\_rename)

and there's a do\_open no there isn't, there's filp\_open.

and a do\_pread64 no there isn't i had to cut/paste sys\_pread64 which was okay because it's pretty basic, just call vfs\_read.

and a do\_mkdir no there isn't so i had to cut/paste that.

basically what i am doing is writing a file system "proxy" module which re-calls back into the filesystem with a prefix onto the front of the pathname.

> > [p.s. i found asm/unistd.h, i found the macros syscall012345  
> > etc., i believe i don't quite understand what these are for, and  
> > may be on the wrong track.]  
>

Linux-Kernel: Re: how do you call userspace syscalls (e.g. sys\_rename) from inside kernel

- > *These are available for you to make syscalls from user mode*
- > *without library support (usually that brand new syscall you added).*
- > *They are basically wrappers that expand into C code. \_syscallx,*
- > *where x is the number of arguments the syscall needs.*

so, it's for use the other way round. okay, thanks for keeping me off a broken line of enquiry.

[oh, and i'll be abandoning this line of enquiry \_entirely\_ if i find that supermount-ng can do the same job – namely manage to keep userspace programs happy when users rip out media]

--  
--

Truth, honesty and respect are rare commodities that all spring from the same well: Love. If you love yourself and everyone and everything around you, funnily and coincidentally enough, life gets a lot better.

--

<http://lkcl.net> lkcl.net </a> <br />

<mailto:lkcl@lkcl.net> lkcl@lkcl.net </a> <br />

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>