Subject: Re: [RFC PATCH 5/5] syscall: sys_fbind() introduced
Posted by hpa on Wed, 15 Aug 2012 16:30:37 GMT
View Forum Message <> Reply to Message

On 08/15/2012 09:22 AM, Stanislav Kinsbursky wrote:
> This syscall allows to bind socket to specified file descriptor.
> Descriptor can be gained by simple open with O_PATH flag.
> Socket node can be created by sys_mknod().
>
> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
> ---
> arch/x86/syscalls/syscall_32.tbl | 1 +
> arch/x86/syscalls/syscall_64.tbl | 1 +
> include/linux/syscalls.h | 1 +
> kernel/sys_ni.c | 3 +++
> net/socket.c | 25 +++++++++++++++++++++++++
> 5 files changed, 31 insertions(+), 0 deletions(-)
>
> diff --git a/arch/x86/syscalls/syscall_32.tbl b/arch/x86/syscalls/syscall_32.tbl
> index 7a35a6e..9594b82 100644
> --- a/arch/x86/syscalls/syscall_32.tbl
> +++ b/arch/x86/syscalls/syscall_32.tbl
> @@ -356,3 +356,4 @@
> 347 i386 process_vm_readv sys_process_vm_readv compat_sys_process_vm_readv
> 348 i386 process_vm_writev sys_process_vm_writev compat_sys_process_vm_writev
> 349 i386 kcmp sys_kcmp
> +350 i386 fbind sys_fbind

i386 uses socketcalls... perhaps it shouldn't (socketcalls are pretty
much an abomination), but for socketcall-based architectures this really
should be a socketcall.

Don't you also need fconnect()? Or is that simply handled by allowing
open() without O_PATH?

 -hpa

--
H. Peter Anvin, Intel Open Source Technology Center
I work for Intel. I don't speak on their behalf.