

---

Subject: Re: [RFC PATCH 0/5] net: socket bind to file descriptor introduced  
Posted by [hpa](#) on Wed, 15 Aug 2012 17:54:15 GMT

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

---

On 08/15/2012 09:52 AM, Ben Pfaff wrote:

> Stanislav Kinsbursky <[skinsbursky@parallels.com](mailto:skinsbursky@parallels.com)> writes:

>

>> This system call is especially required for UNIX sockets, which has name  
>> lenght limitation.

>

> The worst of the name length limitations can be worked around by  
> opening the directory where the socket is to go as a file  
> descriptor, then using /proc/self/fd/<fd>/<basename> as the name  
> of the socket. This technique also works with "connect" and in  
> other contexts where a struct sockaddr is needed. At first  
> glance, it looks like your patches only help with "bind".

>

The really hard part is what to do with things that are supposed to  
return a struct sockaddr. I also have some reservations about using a  
new system call to deal with what at least theoretically is only part of  
one socket domain.

-hpa

---