

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

- > This patch set introduces new socket operation and new system call:
- > sys\_fbind(), which allows to bind socket to opened file.
- > File to bind to can be created by sys\_mknod(S\_IFSOCK) and opened by
- > open(O\_PATH).
- >
- > This system call is especially required for UNIX sockets, which has name
- > lenght limitation.
- >
- > The following series implements...

Thinking about this a little more I have serious reservations about this approach.

Today you are not allowed to bind to an address unless mknod for that file succeeds. Your patch totally changes those semantics.

Name length limitation does not seem to justify this at all.

It is possible today to trivially change into a directory and bind or connect to what would be a long absolute path.

There is also the trick of getting a shorter directory name using /proc/self/fd if you are threaded and can't change the directory.

The obvious choices at this point are

- Teach bind and connect and af\_unix sockets to take longer AF\_UNIX socket path names.
- introduce sockaddr\_fd that can be applied to AF\_UNIX sockets, and teach unix\_bind and unix\_connect how to deal with a second type of sockaddr.  
struct sockaddr\_fd { short fd\_family; short pad; int fd; };
- introduce sockaddr\_unix\_at that takes a directory file descriptor as well as a unix path, and teach unix\_bind and unix\_connect to deal with a second sockaddr type.  
struct sockaddr\_unix\_at { short family; short pad; int dfd; char path[102]; }  
AF\_UNIX\_AT

I don't know what the implications of for breaking connect up into 3 system calls and changing the semantics are and I would really rather not have to think about it.

But it certainly does not look to me like you introduce new systems calls to do what you want.

Eric

---