Subject: Re: [RFC PATCH 0/2] net: connect to UNIX sockets from specified root
Posted by Alan Cox on Fri, 10 Aug 2012 19:24:44 GMT

On Fri, 10 Aug 2012 15:11:50 -0400
"J. Bruce Fields" <bfields@fieldses.org> wrote:

> On Fri, Aug 10, 2012 at 07:26:28PM +0100, Alan Cox wrote:
> > > On that whole subject...
> > >
> > > Do we need a Unix domain socket equivalent to openat()?
> >
> > I don't think so. The name is just a file system indexing trick, it's not
> > really the socket proper. It's little more than "ascii string with
> > permissions attached"
>
> That's overstating the case. As I understand it the address is resolved
> by a pathname lookup like any other--it can follow symlinks, is relative
> to the current working directory and filesystem namespace, etc.

Explicitly for Linux yes - this is not generally true of the AF_UNIX
socket domain and even the permissions aspect isn't guaranteed to be
supported on some BSD environments !

The name is however just a proxy for the socket itself. You don't even
get a device node in the usual sense or the same inode in the file system
space.

Alan