
Subject: Re: [RFC PATCH 0/2] net: connect to UNIX sockets from specified root
Posted by [Stanislav Kinsbursky](#) on Fri, 10 Aug 2012 18:50:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

> On 08/10/2012 05:57 AM, Stanislav Kinsbursky wrote:
>> Today, there is a problem in connecting of local SUNRPC transports. These
>> transports uses UNIX sockets and connection itself is done by rpciod
>> workqueue.
>> But UNIX sockets lookup is done in context of process file system root. I.e.
>> all local transports are connecting in rpciod context.
>> This works nice until we will try to mount NFS from process with other root -
>> for example in container. This container can have it's own (nested) root and
>> rpcbind process, listening on it's own unix sockets. But NFS mount attempt in
>> this container will register new service (Lockd for example) in global rpcbind
>> - not containers's one.
>>
>> This patch set introduces kernel connect helper for UNIX stream sockets and
>> modifies `unix_find_other()` to be able to search from specified root.
>> It also replaces generic socket connect call for local transports by new
>> helper in SUNRPC layer.
>>
>> The following series implements...
> On that whole subject...
>
> Do we need a Unix domain socket equivalent to `openat()`?

It looks like `sys_connectat()` and `sys_bindat()` could be an organic
part of `openat()` and friends family.

But currently I don't have any usage example for them in hands. And the
main problem here, that this syscalls can be used only for unix sockets.
