

> On Fri, Oct 28, 2011 at 01:24:45PM +0400, Stanislav Kinsbursky wrote:  
>> This patch-set was created before you've sent your NFSd plan and we  
>> disacussed Lockd per netns.  
>> So, this sentence: "NFSd service will be per netns too from my pow"  
>> is obsolete. And Lockd will be one for all.  
>  
> I believe lockd should be pert-netns--at least that's what the server  
> needs.  
>  
> (The single lockd thread may handle requests from all netns, but it  
> should behave like a different service depending on netns, so its data  
> structures, etc. will need to be per-ns.  
>

Sure. Looks like we have misunderstanding here. When I said, that Lockd should be one for all, I meaned, that we will have only one kthread for all ns (not one per ns). Private data will be per net ns, of course.

BTW, Bruce, please, have a brief look at my e-mail to linux-nfs@vger.kernel.org named "SUNRPC: non-exclusive pipe creation".  
I've done a lot in "RPC pipefs per net ns" task, and going to send first patches soon. But right now I'm really confused will this non-exclusive pipes creation and almost ready so remove this functionality. But I'm afraid, that I've missed something. Would be greatly appreciate for your opinion about my question.

> --b.  
>  
>> Or you are asking about something else?  
>>  
>>> --b.  
>>>  
>>>> And also we have NFSd file system, which  
>>>> is not virtualized yet.  
>>>>  
>>>> The following series consists of:  
>>>>  
>>>> ---  
>>>>  
>>>> Stanislav Kinsbursky (3):  
>>>> SUNRPC: move rpcbind internals to sunrpc part of network namespace context  
>>>> SUNRPC: optimize net\_ns dereferencing in rpcbind creation calls  
>>>> SUNRPC: optimize net\_ns dereferencing in rpcbind registering calls  
>>>>

```
>>>>
>>>> net/sunrpc/netns.h | 5 ++
>>>> net/sunrpc/rpcb_clnt.c | 103 ++++++-----
>>>> 2 files changed, 61 insertions(+), 47 deletions(-)
>>>>
>>>> --
>>>> Signature
>>
>>
>> --
>> Best regards,
>> Stanislav Kinsbursky
```

```
--
Best regards,
Stanislav Kinsbursky
```

---