

> On Thu, Oct 27, 2011 at 10:10:43PM +0300, Stanislav Kinsbursky wrote:
>> v3:
>> 1) First two patches from previous version were squashed.
>>
>> This patch-set was created in context of clone of git branch:
>> [git://git.linux-nfs.org/projects/trondmy/nfs-2.6.git](http://git.linux-nfs.org/projects/trondmy/nfs-2.6.git)
>> and rebased on tag "v3.1".
>>
>> This patch-set virtualizes rpcbind clients per network namespace context. IOW,
>> each network namespace will have its own pair of rpcbind clients (if they would
>> be created by request).
>>
>> Note:
>> 1) this patch-set depends on "SUNRPC: make rpcbind clients allocated and
>> destroyed dynamically" patch-set which has been send earlier.
>> 2) init_net pointer is still used instead of current->nsproxy->net_ns,
>> because I'm not sure yet about how to virtualize services. I.e. NFS callback
>> services will be per netns. NFSd service will be per netns too from my pow. But
>> Lockd can be per netns or one for all.
>
> I'm not sure what you mean by that; could you explain?
>

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So, this sentence: "NFSd service will be per netns too from my pow" is obsolete.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.
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
