

---

Subject: Re: [PATCH 2/4] NFS: replace per-net client lock by mutex  
Posted by [Myklebust, Trond](#) on Mon, 27 Feb 2012 15:00:18 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Mon, 2012-02-27 at 17:49 +0400, Stanislav Kinsbursky wrote:  
> Lockdep is sad otherwise, because inode mutex is taken on PipeFS dentry  
> creation, which can be called on mount notification, where this per-net client  
> lock is taken on clients list walk.  
>  
> Note: I used simple mutex instead of rw semaphore because of  
> nfs\_put\_client->atomic\_dec\_and\_mutex\_lock() call. Probably, there is a better  
> solution here.  
>  
> Signed-off-by: Stanislav Kinsbursky <[skinsbursky@parallels.com](mailto:skinsbursky@parallels.com)>  
>

This is overkill... We end up blocking NFSv4 callbacks while the  
rpc\_pipefs notifier runs through the nfs\_clients creating or destroying  
idmapper dentries.

Surely the rpc\_pipefs\_event() can take a reference to the nfs\_client and  
then drop the spin\_lock if it sees that it needs to create or destroy a  
dentry?

--  
Trond Myklebust  
Linux NFS client maintainer

NetApp  
[Trond.Myklebust@netapp.com](mailto:Trond.Myklebust@netapp.com)  
[www.netapp.com](http://www.netapp.com)

---