
Subject: Re: [PATCH 2/4] NFS: replace per-net client lock by mutex
Posted by [Stanislav Kinsbursky](#) on Mon, 27 Feb 2012 15:42:49 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>
>>
>
> 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?
>

Sure, thanks for notice.
Looks like this logic also works to SUNRPC clients.
I'll send updated version soon.

--
Best regards,
Stanislav Kinsbursky
