
Subject: Re: [PATCH 4/5] NFS: remove RPC PipeFS mount point reference from blocklayout routines

Posted by [Boaz Harrosh](#) on Mon, 28 May 2012 11:43:53 GMT

[View Forum Message](#) < > [Reply to Message](#)

On 11/29/2011 07:30 PM, Peng Tao wrote:

```
> On Wed, Nov 30, 2011 at 1:19 AM, Trond Myklebust
> <Trond.Myklebust@netapp.com> wrote:
>> On Tue, 2011-11-29 at 11:42 -0500, J. Bruce Fields wrote:
>>> On Tue, Nov 29, 2011 at 11:40:30AM -0500, Trond Myklebust wrote:
>>>> I mean that I'm perfectly entitled to do
>>>>
>>>> 'modprobe -r blocklayoutdriver'
>>>>
>>>> and when I do that, then I expect blkmapd to close the rpc pipe and wait
>>>> for a new one to be created just like rpc.idmapd and rpc.gssd do when I
>>>> remove the nfs and sunrpc modules.
>>>
>>> The rpc pipefs mount doesn't hold a reference on the sunrpc module?
>>
>> I stand corrected: the mount does hold a reference to the sunrpc
>> module.
>> However nothing holds a reference to the blocklayoutdriver module, so
>> the main point that the "blocklayout" pipe can disappear from underneath
>> the blkmapd stands.
> Thanks for the explanation and I agree it can cause problem if user
> reload blocklayout module. I will look into a fix to blkmapd.
>
```

You might want to consider converting to call_usermodehelper()

I know that it greatly simplified our code both in Kernel and in user-mode. And it made nfs-utils maintainer much happier as well.

The speed is not Cardinal here I think. Like in objects it's done once per new device_id

```
> Best,
> Tao
```

Just my \$0.017
Boaz
