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

Posted by [Myklebust, Trond](#) on Tue, 29 Nov 2011 16:40:30 GMT

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

On Tue, 2011-11-29 at 23:30 +0800, Peng Tao wrote:

> On Tue, Nov 29, 2011 at 11:18 PM, Trond Myklebust

> <Trond.Myklebust@netapp.com> wrote:

> > On Tue, 2011-11-29 at 23:10 +0800, Peng Tao wrote:

> > > On Tue, Nov 29, 2011 at 9:35 PM, Myklebust, Trond

> > > <Trond.Myklebust@netapp.com> wrote:

> > > > -----Original Message-----

> > > > From: tao.peng@emc.com [mailto:tao.peng@emc.com]

> > > > Sent: Tuesday, November 29, 2011 7:40 AM

> > > > To: skinsbursky@parallels.com

> > > > Cc: Myklebust, Trond; linux-nfs@vger.kernel.org; xemul@parallels.com;

> > > > neilb@suse.de; netdev@vger.kernel.org; linux-kernel@vger.kernel.org;

> > > > jbottomley@parallels.com; bfields@fieldses.org; davem@davemloft.net;

> > > > devel@openvz.org

> > > > Subject: RE: [PATCH 4/5] NFS: remove RPC PipeFS mount point reference

> > > > from blocklayout routines

> > > >

> > > > > -----Original Message-----

> > > > > From: Stanislav Kinsbursky [mailto:skinsbursky@parallels.com]

> > > > > Sent: Tuesday, November 29, 2011 8:19 PM

> > > > > To: Peng, Tao

> > > > > Cc: Trond.Myklebust@netapp.com; linux-nfs@vger.kernel.org; Pavel

> > > > > Emelianov; neilb@suse.de; netdev@vger.kernel.org;

> > > > > linux-kernel@vger.kernel.org; James Bottomley; bfields@fieldses.org;

> > > > > davem@davemloft.net; devel@openvz.org

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

> > > > > from blocklayout routines

> > > > >

> > > > > > -----Original Message-----

> > > > > > From: linux-nfs-owner@vger.kernel.org

> > > > > > [mailto:linux-nfs-owner@vger.kernel.org] On Behalf Of

> > > > > > Stanislav

> > > > > > Kinsbursky

> > > > > > Sent: Tuesday, November 29, 2011 6:11 PM

> > > > > > To: Trond.Myklebust@netapp.com

> > > > > > Cc: linux-nfs@vger.kernel.org; xemul@parallels.com; neilb@suse.de;

> > > > > > netdev@vger.kernel.org; linux-kernel@vger.kernel.org;

> > > > > > jbottomley@parallels.com; bfields@fieldses.org;

> > > > > > davem@davemloft.net; devel@openvz.org

> > > > > > Subject: [PATCH 4/5] NFS: remove RPC PipeFS mount point reference

> > > > > > from blocklayout routines

> > > > > >

> > > > > This is a cleanup patch. We don't need this reference anymore,
> > > > > because blocklayout pipes dentries now creates and destroys in
> > > > > per-net operations and on PipeFS mount/umount notification.
> > > > > Note that nfs4blocklayout_register_net() now returns 0 instead of
> > > > > -ENOENT in case of PipeFS superblock absence. This is ok, because
> > > > > blocklayout pipe dentry will be created on PipeFS mount event.
> > > > > When is the "pipefs mount event" going to happen? When inserting
> > > > > kernel modules or when user issues
> > > > > mount command?
> > > > >
> > > > >
> > > > > When user issues mount command.
> > > > > Kernel mounts of PipeFS has been removed with all these patch sets
> > > > > I've sent already.
> > > > > Then it is going to break blocklayout user space program blkmapd, which is
> > > > > started before mounting any file system and it tries to open the pipe file
> > > > > when started.
> > > > >
> > > > > Why on earth is blkmapd doing this instead of listening for file creation notifications like the
> > > > > other rpc_pipefs daemons do?
> > > > > Not sure how the original implementer chose this but I think it is
> > > > > likely because we do not expect the pipe file to be created or deleted
> > > > > dynamically.
> > > > >
> > > > > Unless blkmapd can pin the sunrpc module (which it shouldn't be able to)
> > > > > then that assumption would be wrong. Please look into fixing blkmapd...
> > > > > Sorry, I don't quite get it. Do you mean sunrpc module may be removed
> > > > > while nfs/blocklayout modules are still in use? Please explain it a
> > > > > bit. Thanks.

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.

--

Trond Myklebust
Linux NFS client maintainer

NetApp
Trond.Myklebust@netapp.com
www.netapp.com