
Subject: Re: [PATCH 0/5] NFS: create blocklayout pipe per network namespace context

Posted by [Myklebust, Trond](#) on Thu, 05 Jan 2012 20:58:31 GMT

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

On Fri, 2011-12-30 at 17:55 -0500, Trond Myklebust wrote:

> On Tue, 2011-11-29 at 13:10 +0300, Stanislav Kinsbursky wrote:

> > This patch set was created in context of clone of git

> > branch: git://git.linux-nfs.org/projects/trondmy/nfs-2.6.git.

> > tag: v3.1

> >

> > This patch set depends on previous patch sets titled:

> > 1) "SUNRPC: initial part of making pipefs work in net ns"

> > 2) "SUNRPC: cleanup PipeFS for network-namespace-aware users"

> > 3) "SUNRPC: make RPC clients use network-namespace-aware PipeFS routines"

> > 4) "NFS: create clients and IDMAP pipes per network namespace"

> >

> > This patch set is the final part of making SUNRPC PipeFS and it's users work in

> > network namespace context.

> >

> > The following series consists of:

> >

> > ---

> >

> > Stanislav Kinsbursky (5):

> > NFS: handle blocklayout pipe PipeFS dentry by network namespace aware routines

> > NFS: blocklayout pipe creation per network namespace context introduced

> > NFS: blocklayout PipeFS notifier introduced

> > NFS: remove RPC PipeFS mount point reference from blocklayout routines

> > SUNRPC: kernel PipeFS mount point creation routines removed

> >

> >

> > fs/nfs/blocklayout/blocklayout.c | 154 ++++++-----

> > fs/nfs/blocklayout/blocklayout.h | 3 -

> > fs/nfs/blocklayout/blocklayoutdev.c | 5 +

> > fs/nfs/blocklayout/blocklayoutdm.c | 7 +-

> > fs/nfs/inode.c | 1

> > fs/nfs/netns.h | 1

> > include/linux/sunrpc/rpc_pipe_fs.h | 2

> > net/sunrpc/rpc_pipe.c | 21 -----

> > 8 files changed, 137 insertions(+), 57 deletions(-)

>

>

> These patches need rebasing in order to apply on top of 3.2-rc7...

OK. Further testing seems to indicate that we're going to have to postpone merging these patches until the 3.4 merge window.

The problems are twofold:

As the patches stand now in the linux-next tree, they can (and occasionally do) Oops on unmount. The reason was that I rejected the PipeFS notifier patch for the idmapper (due to the reported problem of `nfs_idmap_init/nfs_idmap_quit` being undefined when `CONFIG_NFS_V4` is undefined), and the fact that it is missing causes the unmount at the end of our tests to hit the `BUG()` in `fs/dcache.c:905`. This suggests that we will have the same problem with the pNFS block layout driver, since I still haven't received a rebased update of the 5 'create blocklayout pipe per network namespace context' patches.

The second problem that was highlighted was the fact that as they stand today, these patchsets do not allow for bisection. When we hit the Oops, I had Bryan try to bisect where the problem arose. He ended up pointing at the patch "SUNRPC: handle RPC client pipefs dentries by network namespace aware routine", which is indeed the cause, but which is one of the `_dependencies_` for all the PipeFS notifier patches that fix the problem.

--

Trond Myklebust
Linux NFS client maintainer

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