
Subject: [PATCH v2 4/5] NFS: remove RPC PipeFS mount point reference from blocklayout routines

Posted by [Stanislav Kinsbursky](#) on Tue, 10 Jan 2012 13:04:40 GMT

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

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.

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

fs/nfs/blocklayout/blocklayout.c | 9 +-----
1 files changed, 1 insertions(+), 8 deletions(-)

diff --git a/fs/nfs/blocklayout/blocklayout.c b/fs/nfs/blocklayout/blocklayout.c

index 322d920..2c70202 100644

--- a/fs/nfs/blocklayout/blocklayout.c

+++ b/fs/nfs/blocklayout/blocklayout.c

@@ -1034,7 +1034,7 @@ static struct dentry *nfs4blocklayout_register_net(struct net *net,

pipefs_sb = rpc_get_sb_net(net);
if (!pipefs_sb)

- return ERR_PTR(-ENOENT);

+ return 0;

dentry = nfs4blocklayout_register_sb(pipefs_sb, pipe);

rpc_put_sb_net(net);

return dentry;

@@ -1085,7 +1085,6 @@ static struct pernet_operations nfs4blocklayout_net_ops = {

static int __init nfs4blocklayout_init(void)

{

- struct vfsmount *mnt;

int ret;

dprintk("%s: NFSv4 Block Layout Driver Registering...\n", __func__);

@@ -1095,12 +1094,6 @@ static int __init nfs4blocklayout_init(void)

goto out;

init_waitqueue_head(&bl_wq);

-

- mnt = rpc_get_mount();

- if (IS_ERR(mnt)) {

- ret = PTR_ERR(mnt);

- goto out_remove;

```
- }
ret = rpc_pipefs_notifier_register(&nfs4blocklayout_block);
if (ret)
    goto out_remove;
```

Subject: Re: [PATCH v2 4/5] NFS: remove RPC PipeFS mount point reference from blocklayout routines

Posted by [Myklebust, Trond](#) on Wed, 11 Jan 2012 18:33:13 GMT

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

On Tue, 2012-01-10 at 17:04 +0400, Stanislav Kinsbursky wrote:

```
> 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.
```

```
>
> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
```

```
>
> ---
> fs/nfs/blocklayout/blocklayout.c | 9 +-----
> 1 files changed, 1 insertions(+), 8 deletions(-)
>
> diff --git a/fs/nfs/blocklayout/blocklayout.c b/fs/nfs/blocklayout/blocklayout.c
> index 322d920..2c70202 100644
> --- a/fs/nfs/blocklayout/blocklayout.c
> +++ b/fs/nfs/blocklayout/blocklayout.c
> @@ -1034,7 +1034,7 @@ static struct dentry *nfs4blocklayout_register_net(struct net *net,
>
> pipefs_sb = rpc_get_sb_net(net);
> if (!pipefs_sb)
> - return ERR_PTR(-ENOENT);
> + return 0;
```

This looks like a typo of some sort. The caller in `nfs4blocklayout_net_init()` checks for `IS_ERR()`, not `NULL`. I'll just remove this hunk...

```
> dentry = nfs4blocklayout_register_sb(pipefs_sb, pipe);
> rpc_put_sb_net(net);
> return dentry;
> @@ -1085,7 +1085,6 @@ static struct pernet_operations nfs4blocklayout_net_ops = {
>
> static int __init nfs4blocklayout_init(void)
> {
> - struct vfsmount *mnt;
```

```
> int ret;
>
> dprintk("%s: NFSv4 Block Layout Driver Registering...\n", __func__);
> @@ -1095,12 +1094,6 @@ static int __init nfs4blocklayout_init(void)
> goto out;
>
> init_waitqueue_head(&bl_wq);
> -
> - mnt = rpc_get_mount();
> - if (IS_ERR(mnt)) {
> - ret = PTR_ERR(mnt);
> - goto out_remove;
> - }
> ret = rpc_pipefs_notifier_register(&nfs4blocklayout_block);
> if (ret)
> goto out_remove;
>
> --
> To unsubscribe from this list: send the line "unsubscribe linux-nfs" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html
```

--
Trond Myklebust
Linux NFS client maintainer

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

Subject: Re: [PATCH v2 4/5] NFS: remove RPC PipeFS mount point reference from
blocklayout routines

Posted by [Myklebust, Trond](#) on Wed, 11 Jan 2012 18:36:17 GMT

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

On Wed, 2012-01-11 at 13:33 -0500, Trond Myklebust wrote:

```
> On Tue, 2012-01-10 at 17:04 +0400, Stanislav Kinsbursky wrote:
> > 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.
> >
> > Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
> >
> > ---
```

```

> > fs/nfs/blocklayout/blocklayout.c | 9 +-----
> > 1 files changed, 1 insertions(+), 8 deletions(-)
> >
> > diff --git a/fs/nfs/blocklayout/blocklayout.c b/fs/nfs/blocklayout/blocklayout.c
> > index 322d920..2c70202 100644
> > --- a/fs/nfs/blocklayout/blocklayout.c
> > +++ b/fs/nfs/blocklayout/blocklayout.c
> > @@ -1034,7 +1034,7 @@ static struct dentry *nfs4blocklayout_register_net(struct net *net,
> >
> > pipefs_sb = rpc_get_sb_net(net);
> > if (!pipefs_sb)
> > - return ERR_PTR(-ENOENT);
> > + return 0;
> >
> This looks like a typo of some sort. The caller in
> nfs4blocklayout_net_init() checks for IS_ERR(), not NULL. I'll just
> remove this hunk...

```

Never mind. I see what's going on now... It should be NULL, and not 0 for coding style reasons, though...

--
Trond Myklebust
Linux NFS client maintainer

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

Subject: Re: [PATCH v2 4/5] NFS: remove RPC PipeFS mount point reference from blocklayout routines

Posted by [Stanislav Kinsbursky](#) on Thu, 12 Jan 2012 08:45:25 GMT

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

```

> On Wed, 2012-01-11 at 13:33 -0500, Trond Myklebust wrote:
>> On Tue, 2012-01-10 at 17:04 +0400, Stanislav Kinsbursky wrote:
>>> 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.
>>>
>>> Signed-off-by: Stanislav Kinsbursky<skinsbursky@parallels.com>
>>>
>>> ---

```

```
>>> fs/nfs/blocklayout/blocklayout.c | 9 +-----
>>> 1 files changed, 1 insertions(+), 8 deletions(-)
>>>
>>> diff --git a/fs/nfs/blocklayout/blocklayout.c b/fs/nfs/blocklayout/blocklayout.c
>>> index 322d920..2c70202 100644
>>> --- a/fs/nfs/blocklayout/blocklayout.c
>>> +++ b/fs/nfs/blocklayout/blocklayout.c
>>> @@ -1034,7 +1034,7 @@ static struct dentry *nfs4blocklayout_register_net(struct net *net,
>>>
>>> pipefs_sb = rpc_get_sb_net(net);
>>> if (!pipefs_sb)
>>> - return ERR_PTR(-ENOENT);
>>> + return 0;
>>>
>> This looks like a typo of some sort. The caller in
>> nfs4blocklayout_net_init() checks for IS_ERR(), not NULL. I'll just
>> remove this hunk...
>
> Never mind. I see what's going on now... It should be NULL, and not 0
> for coding style reasons, though...
>
```

Yep, NULL of course, sorry.

--

Best regards,
Stanislav Kinsbursky
