
Subject: Re: [PATCH 4/4] NFS: make nfs_client_lock per net ns
Posted by Myklebust, Trond on Tue, 07 Feb 2012 13:51:00 GMT

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

On Mon, 2012-01-23 at 17:26 +0000, Stanislav Kinsbursky wrote:

> This patch makes nfs_clients_lock allocated per network namespace. All items it
> protects are already network namespace aware.
>
> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
>
> ---
> fs/nfs/client.c | 51 ++++++-----
> fs/nfs/idmap.c | 4 +-+
> fs/nfs/internal.h | 3 ---
> fs/nfs/netns.h | 1 +
> 4 files changed, 32 insertions(+), 27 deletions(-)
>
> diff --git a/fs/nfs/client.c b/fs/nfs/client.c
> index f51b279..9e11d29 100644
> --- a/fs/nfs/client.c
> +++ b/fs/nfs/client.c
> @@ -551,7 +553,7 @@ install_client:
> * - make sure it's ready before returning
> */
> found_client:
> - spin_unlock(&nfs_client_lock);
> + spin_unlock(&nn->nfs_client_lock);
>
> if (new)
> nfs_free_client(new);
> @@ -1041,24 +1043,25 @@ static void nfs_server_insert_lists(struct nfs_server *server)
> struct nfs_client *clp = server->nfs_client;
> struct nfs_net *nn = net_generic(clp->net, nfs_net_id);
>
> - spin_lock(&nfs_client_lock);
> + spin_lock(&nn->nfs_client_lock);
> list_add_tail_rcu(&server->client_link, &clp->cl_superblocks);
> list_add_tail(&server->master_link, &nn->nfs_volume_list);
> clear_bit(NFS_CS_STOP_RENEW, &clp->cl_res_state);
> - spin_unlock(&nfs_client_lock);
> + spin_unlock(&nn->nfs_client_lock);
>
> }
>
> static void nfs_server_remove_lists(struct nfs_server *server)
> {
> struct nfs_client *clp = server->nfs_client;
> + struct nfs_net *nn = net_generic(clp->net, nfs_net_id);

```

>
> - spin_lock(&nfs_client_lock);
> + spin_lock(&nn->nfs_client_lock);
> list_del_rcu(&server->client_link);
> if (clp && list_empty(&clp->cl_superblocks))
>   set_bit(NFS_CS_STOP_RENEW, &clp->cl_res_state);
> list_del(&server->master_link);
> - spin_unlock(&nfs_client_lock);
> + spin_unlock(&nn->nfs_client_lock);
>
> synchronize_rcu();
> }

```

This hunk causes an Oops when nfs_server_remove_lists gets called from nfs4_create_server(). I've applied the following patch to fix it up.

Cheers

Trond

8< -----

>From 5a489156da4fd15dd143f2b21dd9657b97dcef88 Mon Sep 17 00:00:00 2001
 From: Trond Myklebust <Trond.Myklebust@netapp.com>
 Date: Tue, 7 Feb 2012 00:05:11 -0500
 Subject: [PATCH] NFS: Initialise the nfs_net->nfs_client_lock

Ensure that we initialise the nfs_net->nfs_client_lock spinlock.
 Also ensure that nfs_server_remove_lists() doesn't try to
 dereference server->nfs_client before that is initialised.

Signed-off-by: Trond Myklebust <Trond.Myklebust@netapp.com>
 Cc: Stanislav Kinsbursky <skinsbursky@parallels.com>

fs/nfs/client.c | 6 ++++++
 1 files changed, 5 insertions(+), 1 deletions(-)

```

diff --git a/fs/nfs/client.c b/fs/nfs/client.c
index 1a5cd49..f0dacad 100644
--- a/fs/nfs/client.c
+++ b/fs/nfs/client.c
@@ -1055,8 +1055,11 @@ static void nfs_server_insert_lists(struct nfs_server *server)
 static void nfs_server_remove_lists(struct nfs_server *server)
 {
 struct nfs_client *clp = server->nfs_client;
- struct nfs_net *nn = net_generic(clp->net, nfs_net_id);
+ struct nfs_net *nn;

+ if (clp == NULL)
+   return;
+ nn = net_generic(clp->net, nfs_net_id);

```

```
spin_lock(&nn->nfs_client_lock);
list_del_rcu(&server->client_link);
if (clp && list_empty(&clp->cl_superblocks))
@@ -1777,6 +1780,7 @@ void nfs_clients_init(struct net *net)
#endif CONFIG_NFS_V4
idr_init(&nn->cb_ident_idr);
#endif
+ spin_lock_init(&nn->nfs_client_lock);
}

#ifndef CONFIG_PROC_FS
```

--
1.7.7.6

--
Trond Myklebust
Linux NFS client maintainer

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