
Subject: Re: [PATCH] NFS: init client before declaration
Posted by [Chuck Lever](#) on Tue, 22 May 2012 13:47:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

On May 22, 2012, at 8:40 AM, Stanislav Kinsbursky wrote:

```
> Client have to be initialized prior to adding it to per-net clients list,  
> because otherwise there are races, shown below:  
>  
> CPU#0    CPU#1  
> _____  
>  
> nfs_get_client  
> nfs_alloc_client  
> list_add(..., nfs_client_list)  
>     rpc_fill_super  
>     rpc_pipefs_event  
>     nfs_get_client_for_event  
>     __rpc_pipefs_event  
>     (clp->cl_rpcclient is uninitialized)  
>     BUG()  
> init_client  
> clp->cl_rpcclient = ...  
>  
>  
> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
```

This patch collides pretty hard with the server trunking detection work. If you agree this needs to be fixed, the best thing we can do, I guess, is take this patch and drop patch 11, 12, and 13 from my recent patch set, and I'll try to rework for 3.6.

```
> ---  
> fs/nfs/client.c | 22 ++++++-----  
> 1 files changed, 12 insertions(+), 10 deletions(-)  
>  
> diff --git a/fs/nfs/client.c b/fs/nfs/client.c  
> index ae29d4f..9bf4702 100644  
> --- a/fs/nfs/client.c  
> +++ b/fs/nfs/client.c  
> @@ -525,7 +525,7 @@ nfs_get_client(const struct nfs_client_initdata *cl_init,  
>   cl_init->hostname ?: "", cl_init->rpc_ops->version);  
>  
> /* see if the client already exists */  
> - do {  
> + while (1) {  
>   spin_lock(&nn->nfs_client_lock);  
>  
>   clp = nfs_match_client(cl_init);
```

```

> @@ -537,10 +537,18 @@ nfs_get_client(const struct nfs_client_initdata *cl_init,
>   spin_unlock(&nn->nfs_client_lock);
>
>   new = nfs_alloc_client(cl_init);
> - } while (!IS_ERR(new));
> + if (IS_ERR(new)) {
> +   dprintk("--> nfs_get_client() = %ld [failed]\n", PTR_ERR(new));
> +   return new;
> + }
>
> - dprintk("--> nfs_get_client() = %ld [failed]\n", PTR_ERR(new));
> - return new;
> + error = cl_init->rpc_ops->init_client(new, timeparms, ip_addr,
> +     authflavour, noresvport);
> + if (error < 0) {
> +   nfs_put_client(new);
> +   return ERR_PTR(error);
> + }
> +
> /* install a new client and return with it unready */
> install_client:
> @@ -548,12 +556,6 @@ install_client:
>   list_add(&clp->cl_share_link, &nn->nfs_client_list);
>   spin_unlock(&nn->nfs_client_lock);
>
> - error = cl_init->rpc_ops->init_client(clp, timeparms, ip_addr,
> -     authflavour, noresvport);
> - if (error < 0) {
> -   nfs_put_client(clp);
> -   return ERR_PTR(error);
> - }
> - dprintk("--> nfs_get_client() = %p [new]\n", clp);
> - return clp;
>
>
> --
> 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

```

--
Chuck Lever
chuck[dot]lever[at]oracle[dot]com
