
Subject: Re: [PATCH v2] SUNRPC: skip dead but not buried clients on PipeFS events

Posted by [bfields](#) on Wed, 25 Apr 2012 17:30:05 GMT

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

On Fri, Apr 20, 2012 at 06:11:02PM +0400, Stanislav Kinsbursky wrote:

> v2: atomic_inc_return() was replaced by atomic_inc_not_zero().

>

> These clients can't be safely dereferenced if their counter is 0.

I'm pretty confused by how these notifiers work....

rpc_release_client decrements cl_count to zero temporarily, to have it immediately re-incremented by rpc_free_auth.

So if we're called concurrently with rpc_release_client then it's sort of random whether someone gets this callback.

Is that a problem?

Also, is this an existing bug? (In which case Trond should take it now.)

--b.

>
> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
>
> ---
> net/sunrpc/clnt.c | 3 +-+
> 1 files changed, 2 insertions(+), 1 deletions(-)
>
> diff --git a/net/sunrpc/clnt.c b/net/sunrpc/clnt.c
> index 6797246..d10ebc4 100644
> --- a/net/sunrpc/clnt.c
> +++ b/net/sunrpc/clnt.c
> @@ -218,7 +218,8 @@ static struct rpc_clnt *rpc_get_client_for_event(struct net *net, int
event)
> if ((event == RPC_PIPEFS_MOUNT) && clnt->cl_dentry) ||
> ((event == RPC_PIPEFS_UNMOUNT) && !clnt->cl_dentry))
> continue;
> - atomic_inc(&clnt->cl_count);
> + if (atomic_inc_not_zero(&clnt->cl_count) == 0)
> + continue;
> spin_unlock(&sn->rpc_client_lock);
> return clnt;
> }
>
