Subject: Re: [PATCH 0/3] lockd: use per-net refrence-counted NSM clients
Posted by Chuck Lever on Mon, 17 Sep 2012 15:10:30 GMT
View Forum Message <> Reply to Message

On Sep 17, 2012, at 6:49 AM, Stanislav Kinsbursky wrote:

>>
>> On Sep 14, 2012, at 1:38 PM, Myklebust, Trond wrote:
>>
>>> On Fri, 2012-09-14 at 13:01 -0400, Chuck Lever wrote:
>>>> What happens if statd is restarted?
>>>
>>> Nothing unusual. Why?
>>
>> The NSM upcall transport is a potential application for TCP + softconn, now that a persistent rpc_clnt is used. It just depends on what failure mode we'd like to optimize for.
>>
>
> I don't understand, where the problem is.
> Could you be more specific, please?

I'm suggesting an enhancement.

The change is to use TCP for the NSM upcall transport, and set RPC_TASK_SOFTCONN on the individual RPCs. The advantage of this is that the kernel could discover when statd is not running and fail the upcall immediately, rather than waiting possibly many seconds for each upcall RPC to time out.

The client already has a check in the mount.nfs command to see that statd is running, likely to avoid this lengthly timeout. Since the client already has long-standing logic to avoid it, I think the benefit would be mostly on the server side.

But this change can be done at some later point.

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