

---

Subject: Re: [PATCH 0/3] lockd: use per-net refrence-counted NSM clients

Posted by [Chuck Lever](#) on Fri, 14 Sep 2012 19:10:02 GMT

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

---

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.

>

>> Sent from my iPhonespam SPAMSPAM

>>

>> On Sep 14, 2012, at 10:25 AM, Stanislav Kinsbursky <[skinsbursky@parallels.com](mailto:skinsbursky@parallels.com)> wrote:

>>

>>> This is a bug fix for [https://bugzilla.redhat.com/show\\_bug.cgi?id=830862](https://bugzilla.redhat.com/show_bug.cgi?id=830862).

>>>

>>> The problem is that with NFSv4 mount in container (with separated mount

>>> namespace) and active lock on it, dying child reaped of this container will

>>> try to umount NFS and doing this will try to create RPC client to send

>>> unmonitor request to statd.

>>> But creation of RCP client requires valid current->nsproxy (for operation with

>>> utsname()) and during umount on child reaper exit it's equal to zero.

>>>

>>> Proposed solution is to introduce refrence-counter per-net NSM client, which

>>> is created on fist monitor call and destroyed after the 1st monitor call.

>>>

>>> The following series implements...

>>>

>>> ---

>>>

>>> Stanislav Kinsbursky (3):

>>> lockd: use rpc client's cl\_nodename for id encoding

>>> lockd: per-net NSM client creation and destruction helpers introduced

>>> lockd: create and use per-net NSM RPC clients on MON/UNMON requests

>>>

>>>

>>> fs/lockd/mon.c | 91 ++++++-----

>>> fs/lockd/netns.h | 4 ++

>>> fs/lockd/svc.c | 1 +

>>> 3 files changed, 77 insertions(+), 19 deletions(-)

>>>

>>>

>>> --

>>> 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](http://www.netapp.com)

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

---