Subject: [PATCH v2 0/3] lockd: use per-net refrence-counted NSM clients
Posted by Stanislav Kinsbursky on Tue, 18 Sep 2012 09:37:06 GMT
View Forum Message <> Reply to Message

v2:
1) NSM transport is TCP based now. And all RPC tasks now called with RPC_TASK_SOFTCONN. 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.
2) XDR layer violation (reference to upper RPC client cl_hostname) was replaced by passing the string as a part of nlm_args structure.

This is a bug fix for https://bugzilla.redhat.com/show_bug.cgi?id=830862.

The problem is that with NFSv4 mount in container (with separated mount namesapce) 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 lst monitor call.

The following series implements...

Stanislav Kinsbursky (3):
    lockd: per-net NSM client creation and destruction helpers introduced
    lockd: use rpc client's cl_nodename for id encoding
    lockd: create and use per-net NSM RPC clients on MON/UNMON requests

```
 fs/lockd/mon.c   |   86 +++++++++++++++++++++++++++++++++++++++++++-----------
 fs/lockd/netns.h |    4 +++
 fs/lockd/svc.c   |    1 +
 3 files changed, 74 insertions(+), 17 deletions(-)
```