

> On Wed, Jul 25, 2012 at 04:55:45PM +0400, Stanislav Kinsbursky wrote:  
>> Bruce, I feel this patch set is ready for inclusion.  
>>  
>> v2:  
>> 1) Rebase on Bruce's "for-3.6" branch.  
>>  
>> This patch set makes grace period and hosts reclaiming network namespace  
>> aware.  
>  
> On a quick skim--yes, that looks reasonable to me.  
>  
> It doesn't help with active/active cluster exports, because in that case  
> we need some additional coordination between nfsd's.  
>  
> But it looks good enough to handle the case where each filesystem is  
> exported from at most one server at a time, which is more than we  
> currently handle.  
>  
> It's a little late for 3.6. Also I get the impression Al Viro has some  
> lockd rework in progress, which we may want to wait for.  
>  
> So I'll likely look again into queueing this up for 3.7 once 3.6-rc1 is  
> out.

Ok.  
Will Al Viro's lockd rework be a part of 3.6 kernel?

>  
> --b.  
>  
>>  
>> Main ideas:  
>> 1) moving of  
>>  
>> unsigned long next\_gc;  
>> unsigned long nrhosts;  
>>  
>> struct delayed\_work grace\_period\_end;  
>> struct lock\_manager lockd\_manager;  
>> struct list\_head grace\_list;  
>>

```

>> to per-net Lockd data.
>>
>> 2) moving of
>>
>> struct lock_manager nfsd4_manager;
>>
>> to per-net NFSd data.
>>
>> 3) shutdown + gc of NLM hosts done now network namespace aware.
>>
>> 4) restart_grace() now works only for init_net.
>>
>> The following series implements...
>>
>> ---
>>
>> Stanislav Kinsbursky (15):
>>   LockD: mark host per network namespace on garbage collect
>>   LockD: make garbage collector network namespace aware.
>>   LockD: manage garbage collection timeout per networks namespace
>>   LockD: manage used host count per networks namespace
>>   Lockd: host complaining function introduced
>>   Lockd: add more debug to host shutdown functions
>>   LockD: manage grace period per network namespace
>>   LockD: make lockd manager allocated per network namespace
>>   NFSd: make nfsd4_manager allocated per network namespace context.
>>   SUNRPC: service request network namespace helper introduced
>>   LockD: manage grace list per network namespace
>>   LockD: pass actual network namespace to grace period management functions
>>   Lockd: move grace period management from lockd() to per-net functions
>>   NFSd: make grace end flag per network namespace
>>   NFSd: make boot_time variable per network namespace
>>
>>
>> fs/lockd/grace.c      | 16 ++++++
>> fs/lockd/host.c       | 92 ++++++-----
>> fs/lockd/netns.h      | 7 +++
>> fs/lockd/svc.c        | 43 ++++++-----
>> fs/lockd/svc4proc.c   | 13 +++-
>> fs/lockd/svclock.c    | 16 +++-
>> fs/lockd/svcproc.c    | 15 ++++
>> fs/lockd/svcsubs.c    | 19 +++++-
>> fs/nfs/callback_xdr.c | 4 +-
>> fs/nfsd/export.c      | 4 +-
>> fs/nfsd/netns.h       | 4 ++
>> fs/nfsd/nfs4idmap.c   | 4 +-
>> fs/nfsd/nfs4proc.c    | 18 ++++
>> fs/nfsd/nfs4state.c   | 104 ++++++-----

```

```
>> fs/nfsd/state.h      | 4 +-  
>> include/linux/fs.h    | 5 +-  
>> include/linux/lockd/lockd.h | 6 +-  
>> include/linux/sunrpc/svc.h | 2 +  
>> 18 files changed, 231 insertions(+), 145 deletions(-)  
>>
```

--

Best regards,  
Stanislav Kinsbursky

---