
Subject: [PATCH v2] NFSd: simplify locking in nfsd_recall_delegations()

Posted by [Stanislav Kinsbursky](#) on Tue, 03 Jul 2012 13:02:24 GMT

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

v2: rebased on 3.5-rc3 kernel

Holding of client_mutex looks redundant here - recall_lock looks enough.

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

fs/nfsd/nfs4state.c | 2 --

1 files changed, 0 insertions(+), 2 deletions(-)

diff --git a/fs/nfsd/nfs4state.c b/fs/nfsd/nfs4state.c

index 33eccee..480ecc9 100644

--- a/fs/nfsd/nfs4state.c

+++ b/fs/nfsd/nfs4state.c

@@ -4622,7 +4622,6 @@ void nfsd_recall_delegations(u64 num)

LIST_HEAD(victims);

struct nfs4_delegation *dp, *dnext;

- nfs4_lock_state();

spin_lock(&recall_lock);

count = nfsd_process_n_delegations(num, &victims);

list_for_each_entry_safe(dp, dnext, &victims, dl_recall_lru) {

@@ -4630,7 +4629,6 @@ void nfsd_recall_delegations(u64 num)

nfsd_break_one_deleg(dp);

}

spin_unlock(&recall_lock);

- nfs4_unlock_state();

printk(KERN_INFO "NFSD: Recalled %d delegations", count);

}
