
Subject: [PATCH v2 15/15] NFSd: make boot_time variable per network namespace
Posted by Stanislav Kinsbursky on Wed, 25 Jul 2012 12:57:45 GMT

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

NFSd's boot_time represents grace period start point in time.

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

```
fs/nfsd/netns.h    |  1 +
fs/nfsd/nfs4state.c | 39 ++++++-----+
fs/nfsd/state.h    |  1 +
3 files changed, 25 insertions(+), 16 deletions(-)
```

```
diff --git a/fs/nfsd/netns.h b/fs/nfsd/netns.h
```

```
index b6deebd..65c2431 100644
```

```
--- a/fs/nfsd/netns.h
```

```
+++ b/fs/nfsd/netns.h
```

```
@@ -37,6 +37,7 @@ struct nfsd_net {
```

```
    struct lock_manager nfsd4_manager;
```

```
    bool grace_ended;
```

```
    + time_t boot_time;
```

```
};
```

```
extern int nfsd_net_id;
```

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

```
index 2c53447..d683f32 100644
```

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

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

```
@@ -52,7 +52,6 @@
```

```
/* Globals */
```

```
time_t nfsd4_lease = 90; /* default lease time */
```

```
time_t nfsd4_grace = 90;
```

```
-static time_t boot_time;
```

```
#define all_ones {{~0,~0},~0}
```

```
static const stateid_t one_stateid = {
```

```
@@ -1055,12 +1054,12 @@ renew_client(struct nfs4_client *clp)
```

```
/* SETCLIENTID and SETCLIENTID_CONFIRM Helper functions */
```

```
static int
```

```
-STALE_CLIENTID(clientid_t *clid)
```

```
+STALE_CLIENTID(clientid_t *clid, struct nfsd_net *nn)
```

```
{
```

```
- if (clid->cl_boot == boot_time)
```

```
+ if (clid->cl_boot == nn->boot_time)
```

```
    return 0;
```

```
    dprintk("NFSD stale clientid (%08x/%08x) boot_time %08lx\n",
```

```

- clid->cl_boot, clid->cl_id, boot_time);
+ clid->cl_boot, clid->cl_id, nn->boot_time);
return 1;
}

@@ -1241,8 +1240,9 @@ same_creds(struct svc_cred *cr1, struct svc_cred *cr2)
static void gen_clid(struct nfs4_client *clp)
{
    static u32 current_clientid = 1;
+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

- clp->cl_clientid.cl_boot = boot_time;
+ clp->cl_clientid.cl_boot = nn->boot_time;
    clp->cl_clientid.cl_id = current_clientid++;
}

@@ -2225,8 +2225,9 @@ nfsd4_setclientid_confirm(struct svc_rqst *rqstp,
nfs4_verifier confirm = setclientid_confirm->sc_confirm;
clientid_t * clid = &setclientid_confirm->sc_clientid;
__be32 status;
+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

- if (STALE_CLIENTID(clid))
+ if (STALE_CLIENTID(clid, nn))
    return nfserr_stale_clientid;
nfs4_lock_state();

@@ -2585,8 +2586,9 @@ nfsd4_process_open1(struct nfsd4_compound_state *cstate,
unsigned int strhashval;
struct nfs4_openowner *oo = NULL;
__be32 status;
+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

- if (STALE_CLIENTID(&open->op_clientid))
+ if (STALE_CLIENTID(&open->op_clientid, nn))
    return nfserr_stale_clientid;
/*
 * In case we need it later, after we've already created the
@@ -3094,12 +3096,13 @@ nfsd4_renew(struct svc_rqst *rqstp, struct nfsd4_compound_state
*cstate,
{
    struct nfs4_client *clp;
    __be32 status;
+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

nfs4_lock_state();
dprintk("process_renew(%08x/%08x): starting\n",
    clid->cl_boot, clid->cl_id);

```

```

status = nfserr_stale_clientid;
- if (STALE_CLIENTID(clid))
+ if (STALE_CLIENTID(clid, nn))
    goto out;
    clp = find_confirmed_client(clid);
    status = nfserr_expired;
@@ -3129,7 +3132,7 @@ nfsd4_end_grace(struct net *net)

dprintk("NFSD: end of grace period\n");
nn->grace_ended = true;
- nfsd4_record_grace_done(net, boot_time);
+ nfsd4_record_grace_done(net, nn->boot_time);
locks_end_grace(&nn->nfsd4_manager);
/*
 * Now that every NFSv4 client has had the chance to recover and
@@ -3235,9 +3238,9 @@ static inline __be32 nfs4_check_fh(struct svc_fh *fhp, struct
nfs4_ol_stateid *s
}

static int
-STALE_STATEID(stateid_t *stateid)
+STALE_STATEID(stateid_t *stateid, struct nfsd_net *nn)
{
- if (stateid->si_opaque.so_clid.cl_boot == boot_time)
+ if (stateid->si_opaque.so_clid.cl_boot == nn->boot_time)
    return 0;
dprintk("NFSD: stale stateid " STATEID_FMT "!\n",
    STATEID_VAL(stateid));
@@ -3372,10 +3375,11 @@ static __be32 nfsd4_validate_stateid(struct nfs4_client *cl, stateid_t
*stateid)
static __be32 nfsd4_lookup_stateid(stateid_t *stateid, unsigned char typemask, struct nfs4_stid
**s)
{
    struct nfs4_client *cl;
+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

if (ZERO_STATEID(stateid) || ONE_STATEID(stateid))
    return nfserr_bad_stateid;
- if (STALE_STATEID(stateid))
+ if (STALE_STATEID(stateid, nn))
    return nfserr_stale_stateid;
    cl = find_confirmed_client(&stateid->si_opaque.so_clid);
    if (!cl)
@@ -4047,6 +4051,7 @@ nfsd4_lock(struct svc_rqst *rqstp, struct nfsd4_compound_state
*cstate,
    bool new_state = false;
    int lkflg;
    int err;

```

```

+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

dprintk("NFSD: nfsd4_lock: start=%Ld length=%Ld\n",
       (long long) lock->lk_offset,
@@ -4073,7 +4078,7 @@ nfsd4_lock(struct svc_rqst *rqstp, struct nfsd4_compound_state
*cstate,
       sizeof(clientid_t));

status = nfserr_stale_clientid;
- if (STALE_CLIENTID(&lock->lk_new_clientid))
+ if (STALE_CLIENTID(&lock->lk_new_clientid, nn))
    goto out;

/* validate and update open stateid and open seqid */
@@ -4207,6 +4212,7 @@ nfsd4_lockt(struct svc_rqst *rqstp, struct nfsd4_compound_state
*cstate,
struct file_lock file_lock;
struct nfs4_lockowner *lo;
__be32 status;
+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

if (locks_in_grace(SVC_NET(rqstp)))
    return nfserr_grace;
@@ -4217,7 +4223,7 @@ nfsd4_lockt(struct svc_rqst *rqstp, struct nfsd4_compound_state
*cstate,
nfs4_lock_state();

status = nfserr_stale_clientid;
- if (!nfsd4_has_session(cstate) && STALE_CLIENTID(&lockt->lt_clientid))
+ if (!nfsd4_has_session(cstate) && STALE_CLIENTID(&lockt->lt_clientid, nn))
    goto out;

if ((status = fh_verify(rqstp, &cstate->current_fh, S_IFREG, 0)))
@@ -4366,6 +4372,7 @@ nfsd4_release_lockowner(struct svc_rqst *rqstp,
struct list_head matches;
unsigned int hashval = ownerstr_hashval(clid->cl_id, owner);
__be32 status;
+ struct nfsd_net *nn = net_generic(&init_net, nfsd_net_id);

dprintk("nfsd4_release_lockowner clientid: (%08x/%08x):\n",
       clid->cl_boot, clid->cl_id);
@@ -4373,7 +4380,7 @@ nfsd4_release_lockowner(struct svc_rqst *rqstp,
/* XXX check for lease expiration */

status = nfserr_stale_clientid;
- if (STALE_CLIENTID(clid))
+ if (STALE_CLIENTID(clid, nn))
    return status;

```

```
nfs4_lock_state();
@@ -4700,7 +4707,7 @@ nfs4_state_start(void)
 */
get_net(net);
nfsd4_client_tracking_init(net);
- boot_time = get_seconds();
+ nn->boot_time = get_seconds();
locks_start_grace(net, &nn->nfsd4_manager);
nn->grace-ended = false;
printk(KERN_INFO "NFSD: starting %ld-second grace period\n",
diff --git a/fs/nfsd/state.h b/fs/nfsd/state.h
index 981ef10..e617314 100644
--- a/fs/nfsd/state.h
+++ b/fs/nfsd/state.h
@@ -450,6 +450,7 @@ static inline struct nfs4_ol_stateid *openlockstateid(struct nfs4_stid *s)
#define WR_STATE      0x00000020

struct nfsd4_compound_state;
+struct nfsd_net;

extern __be32 nfs4_preprocess_stateid_op(struct net *net,
struct nfsd4_compound_state *cstate,
```
