
Subject: [PATCH v2 2/6] NFS: make NFS client allocated per network namespace context

Posted by [Stanislav Kinsbursky](#) on Tue, 10 Jan 2012 12:12:46 GMT

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

This patch adds new net variable to nfs_client structure. This variable is set on NFS client creation and checked during matching NFS client search. Initially current->nsproxy->net_ns is used as network namespace owner for new NFS client to create. This network namespace pointer is set during mount options parsing and thus can be passed from user-space utils in future if will be necessary.

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

```
---
fs/nfs/client.c      | 16 ++++++-----
fs/nfs/internal.h   |  1 +
fs/nfs/super.c      |  3 +++
include/linux/nfs_fs_sb.h |  1 +
4 files changed, 18 insertions(+), 3 deletions(-)
```

```
diff --git a/fs/nfs/client.c b/fs/nfs/client.c
```

```
index 873bf00..4fb5cae 100644
```

```
--- a/fs/nfs/client.c
```

```
+++ b/fs/nfs/client.c
```

```
@@ -135,6 +135,7 @@ struct nfs_client_initdata {
    const struct nfs_rpc_ops *rpc_ops;
    int proto;
    u32 minorversion;
+ struct net *net;
};
```

```
/*
@@ -189,6 +190,7 @@ static struct nfs_client *nfs_alloc_client(const struct nfs_client_initdata
*cl_
    if (!IS_ERR(cred))
        clp->cl_machine_cred = cred;
    nfs_fscache_get_client_cookie(clp);
+ clp->net = cl_init->net;
```

```
    return clp;
```

```
@@ -476,6 +478,9 @@ static struct nfs_client *nfs_match_client(const struct nfs_client_initdata
*dat
    /* Match the full socket address */
    if (!nfs_sockaddr_cmp(sap, clap))
        continue;
+ /* Match network namespace */
```

```

+ if (clp->net != data->net)
+ continue;

    atomic_inc(&clp->cl_count);
    return clp;
@@ -826,6 +831,7 @@ static int nfs_init_server(struct nfs_server *server,
    .addrlen = data->nfs_server.addrlen,
    .rpc_ops = &nfs_v2_clientops,
    .proto = data->nfs_server.protocol,
+ .net = data->net,
};
struct rpc_timeout timeparms;
struct nfs_client *clp;
@@ -1387,7 +1393,7 @@ static int nfs4_set_client(struct nfs_server *server,
    const char *ip_addr,
    rpc_authflavor_t authflavour,
    int proto, const struct rpc_timeout *timeparms,
- u32 minorversion)
+ u32 minorversion, struct net *net)
{
    struct nfs_client_initdata cl_init = {
        .hostname = hostname,
@@ -1396,6 +1402,7 @@ static int nfs4_set_client(struct nfs_server *server,
    .rpc_ops = &nfs_v4_clientops,
    .proto = proto,
    .minorversion = minorversion,
+ .net = net,
};
struct nfs_client *clp;
int error;
@@ -1447,6 +1454,7 @@ struct nfs_client *nfs4_set_ds_client(struct nfs_client* mds_clp,
    .rpc_ops = &nfs_v4_clientops,
    .proto = ds_proto,
    .minorversion = mds_clp->cl_minorversion,
+ .net = mds_clp->net,
};
struct rpc_timeout ds_timeout = {
    .to_initval = 15 * HZ,
@@ -1573,7 +1581,8 @@ static int nfs4_init_server(struct nfs_server *server,
    data->auth_flavors[0],
    data->nfs_server.protocol,
    &timeparms,
- data->minorversion);
+ data->minorversion,
+ data->net);
if (error < 0)
    goto error;

```

```

@@ -1670,7 +1679,8 @@ struct nfs_server *nfs4_create_referral_server(struct nfs_clone_mount
*data,
    data->authflavor,
    parent_server->client->cl_xprt->prot,
    parent_server->client->cl_timeout,
-   parent_client->cl_mvops->minor_version);
+   parent_client->cl_mvops->minor_version,
+   parent_client->net);
    if (error < 0)
        goto error;

```

```

diff --git a/fs/nfs/internal.h b/fs/nfs/internal.h
index 3f4d957..c60fa0b 100644

```

```

--- a/fs/nfs/internal.h
+++ b/fs/nfs/internal.h
@@ -123,6 +123,7 @@ struct nfs_parsed_mount_data {
    } nfs_server;

    struct security_mnt_opts lsm_opts;
+   struct net *net;
};

```

```

/* mount_clnt.c */
diff --git a/fs/nfs/super.c b/fs/nfs/super.c
index b31bdf4..7495039 100644

```

```

--- a/fs/nfs/super.c
+++ b/fs/nfs/super.c
@@ -53,6 +53,7 @@
#include <linux/nfs_xdr.h>
#include <linux/magic.h>
#include <linux/parser.h>
+#include <linux/nsproxy.h>

#include <asm/system.h>
#include <asm/uaccess.h>
@@ -1094,6 +1095,8 @@ static int nfs_parse_mount_options(char *raw,

```

```

    free_secdata(secdata);

+   mnt->net = current->nsproxy->net_ns;
+
    while ((p = strsep(&raw, ",")) != NULL) {
        substring_t args[MAX_OPT_ARGS];
        unsigned long option;
diff --git a/include/linux/nfs_fs_sb.h b/include/linux/nfs_fs_sb.h
index b5479df..71fd605 100644
--- a/include/linux/nfs_fs_sb.h
+++ b/include/linux/nfs_fs_sb.h

```

```
@@ -85,6 +85,7 @@ struct nfs_client {
#endif

    struct server_scope *server_scope; /* from exchange_id */
+ struct net *net;
};

/*
```
