
Subject: [PATCH] NFS: decode destination address in proper network namespace context

Posted by [Stanislav Kinsbursky](#) on Thu, 19 Jan 2012 15:05:57 GMT

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

This patch replaces "init_net" with NFS client's owner net in rpc_pton() call in decode_ds_addr().

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

fs/nfs/nfs4filelayoutdev.c | 7 ++++---
1 files changed, 4 insertions(+), 3 deletions(-)

diff --git a/fs/nfs/nfs4filelayoutdev.c b/fs/nfs/nfs4filelayoutdev.c

index 0d8b952..6eb59b0 100644

--- a/fs/nfs/nfs4filelayoutdev.c

+++ b/fs/nfs/nfs4filelayoutdev.c

@@ -378,7 +378,7 @@ out:

* Currently only supports ipv4, ipv6 and one multi-path address.

*/

static struct nfs4_pnfs_ds_addr *

-decode_ds_addr(struct xdr_stream *stream, gfp_t gfp_flags)

+decode_ds_addr(struct net *net, struct xdr_stream *stream, gfp_t gfp_flags)

{

struct nfs4_pnfs_ds_addr *da = NULL;

char *buf, *portstr;

@@ -457,7 +457,7 @@ decode_ds_addr(struct xdr_stream *stream, gfp_t gfp_flags)

INIT_LIST_HEAD(&da->da_node);

- if (!rpc_pton(&init_net, buf, portstr-buf, (struct sockaddr *)&da->da_addr,

+ if (!rpc_pton(net, buf, portstr-buf, (struct sockaddr *)&da->da_addr,

sizeof(da->da_addr))) {

dprintk("%s: error parsing address %s\n", __func__, buf);

goto out_free_da;

@@ -625,7 +625,8 @@ decode_device(struct inode *ino, struct pnfs_device *pdev, gfp_t gfp_flags)

mp_count = be32_to_cpu(p); /* multipath count */

for (j = 0; j < mp_count; j++) {

- da = decode_ds_addr(&stream, gfp_flags);

+ da = decode_ds_addr(NFS_SERVER(ino)->nfs_client->net,

+ &stream, gfp_flags);

if (da)

list_add_tail(&da->da_node, &dsaddrs);

}