

---

Subject: [PATCH 1/2] NFS: replace global bl\_mount\_reply with per-net one  
Posted by Stanislav Kinsbursky on Sun, 11 Mar 2012 14:20:23 GMT

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

---

This global variable is used for blocklayout downcall and thus can be corrupted if case of existence of multiple networks namespaces.

Signed-off-by: Stanislav Kinsbursky <[skinsbursky@parallels.com](mailto:skinsbursky@parallels.com)>

---

```
fs/nfs/blocklayout/blocklayout.h |  5 -----
fs/nfs/blocklayout/blocklayoutdev.c |  9 ++++++-
fs/nfs/netns.h                 |  6 ++++++
3 files changed, 11 insertions(+), 9 deletions(-)
```

```
diff --git a/fs/nfs/blocklayout/blocklayout.h b/fs/nfs/blocklayout/blocklayout.h
index 0966b39..58ac861 100644
--- a/fs/nfs/blocklayout/blocklayout.h
+++ b/fs/nfs/blocklayout/blocklayout.h
@@ -153,11 +153,6 @@ BLK_LSEG2EXT(struct pnfs_layout_segment *lseg)
    return BLK_LO2EXT(lseg->pls_layout);
}
```

```
-struct bl_dev_msg {
- int32_t status;
- uint32_t major, minor;
-};
-
 struct bl_msg_hdr {
 u8 type;
 u16 totallen; /* length of entire message, including hdr itself */
diff --git a/fs/nfs/blocklayout/blocklayoutdev.c b/fs/nfs/blocklayout/blocklayoutdev.c
index b48f782..1d58642 100644
--- a/fs/nfs/blocklayout/blocklayoutdev.c
+++ b/fs/nfs/blocklayout/blocklayoutdev.c
@@ -79,15 +79,16 @@ int nfs4_blkdev_put(struct block_device *bdev)
    return blkdev_put(bdev, FMODE_READ);
}
```

```
-static struct bl_dev_msg bl_mount_reply;
-
 ssize_t bl_pipe_downcall(struct file *filp, const char __user *src,
 size_t mlen)
{
+ struct nfs_net *nn = net_generic(filp->f_dentry->d_sb->s_fs_info,
+     nfs_net_id);
+
 if (mlen != sizeof (struct bl_dev_msg))
```

```

return -EINVAL;

- if (copy_from_user(&bl_mount_reply, src, mlen) != 0)
+ if (copy_from_user(&nn->bl_mount_reply, src, mlen) != 0)
    return -EFAULT;

wake_up(&bl_wq);
@@ -118,10 +119,10 @@ nfs4_blk_decode_device(struct nfs_server *server,
};

uint8_t *dataptr;
DECLARE_WAITQUEUE(wq, current);
- struct bl_dev_msg *reply = &bl_mount_reply;
int offset, len, i, rc;
struct net *net = server->nfs_client->net;
struct nfs_net *nn = net_generic(net, nfs_net_id);
+ struct bl_dev_msg *reply = &nn->bl_mount_reply;

dprintk("%s CREATING PIPEFS MESSAGE\n", __func__);
dprintk("%s: deviceid: %s, mincount: %d\n", __func__, dev->dev_id.data,
diff --git a/fs/nfs/netns.h b/fs/nfs/netns.h
index 7baad89..73425f5 100644
--- a/fs/nfs/netns.h
+++ b/fs/nfs/netns.h
@@ -4,9 +4,15 @@
#include <net/net_namespace.h>
#include <net/netns/generic.h>

+struct bl_dev_msg {
+ int32_t status;
+ uint32_t major, minor;
+};
+
 struct nfs_net {
 struct cache_detail *nfs_dns_resolve;
 struct rpc_pipe *bl_device_pipe;
+ struct bl_dev_msg bl_mount_reply;
 struct list_head nfs_client_list;
 struct list_head nfs_volume_list;
 #ifdef CONFIG_NFS_V4

```

---