Subject: [PATCH 2/7] SUNRPC: hold current network namespace while pipefs superblock is active
Posted by Stanislav Kinsbursky on Fri, 28 Oct 2011 14:26:34 GMT
View Forum Message <> Reply to Message

We want to be sure that network namespace is still alive while we have pipefs
mounted.
This will be required later, when RPC pipefs will be mounting only from
user-space context.

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

---
 net/sunrpc/rpc_pipe.c | 14 ++++++++++++-
 1 files changed, 13 insertions(+), 1 deletions(-)

diff --git a/net/sunrpc/rpc_pipe.c b/net/sunrpc/rpc_pipe.c
index ad517bd..a717564 100644
--- a/net/sunrpc/rpc_pipe.c
+++ b/net/sunrpc/rpc_pipe.c
@@ -27,6 +27,9 @@
 #include <linux/workqueue.h>
 #include <linux/sunrpc/rpc_pipe_fs.h>
 #include <linux/sunrpc/cache.h>
+#include <linux/nsproxy.h>
+
+#include "netns.h"

 static struct vfsmount *rpc_mnt __read_mostly;
 static int rpc_mount_count;
@@ -1024,6 +1027,7 @@ rpc_fill_super(struct super_block *sb, void *data, int silent)
  }
  if (rpc_populate(root, files, RPCAUTH_lockd, RPCAUTH_RootEOF, NULL))
   return -ENOMEM;
+ sb->s_fs_info = get_net(net);
  return 0;
 }

@@ -1034,11 +1038,19 @@ rpc_mount(struct file_system_type *fs_type,
  return mount_ns(fs_type, flags, current->nsproxy->net_ns, rpc_fill_super);
 }

+void rpc_kill_sb(struct super_block *sb)
+{
+ struct net *net = sb->s_fs_info;
+
+ put_net(net);
+ kill_litter_super(sb);

```
+}
+
 static struct file_system_type rpc_pipe_fs_type = {
  .owner  = THIS_MODULE,
  .name  = "rpc_pipefs",
  .mount  = rpc_mount,
- .kill_sb = kill_litter_super,
+ .kill_sb = rpc_kill_sb,
 };

 static void
```