
Subject: [PATCH 13/13] Make all proc entries accessible in a namespace
Posted by [Pavel Emelianov](#) on Thu, 24 May 2007 13:13:16 GMT

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

This is a hack that must not go to any tree, but this makes patch testing from shell possible.

When creating a new pid namespace one need to mount a new proc instance in this space context. But the /proc will contain pids and self symlink only and this breaks some important tools like ps and sshd.

They expect /proc/stat, /proc/mounts and some other entries to be present.

Without this patch the only possible way to test the namespace is to make all the operations performed from self-made programs. With it one may use general command line tools.

This patch makes them visible, but this is ugly. Correct fix must make some kind of proc_entries namespace but this is another interesting task.

Signed-off-by: Pavel Emelianov <xemul@openvz.org>

```
diff --git a/fs/proc/generic.c b/fs/proc/generic.c
index 9401a6f..eec1eb6 100644
--- a/fs/proc/generic.c
+++ b/fs/proc/generic.c
@@ -393,6 +393,10 @@ struct dentry *proc_lookup(struct inode
    spin_lock(&proc_subdir_lock);
    de = PDE(dir);
    if (de) {
+#ifdef CONFIG_PID_NS
+       if (de->parent == de)
+           de = &proc_root;
+#endif
        for (de = de->subdir; de ; de = de->next) {
            if (de->namelen != dentry->d_name.len)
                continue;
@@ -446,6 +450,10 @@ int proc_readdir(struct file * filp,
    ret = -EINVAL;
    goto out;
}
+#ifdef CONFIG_PID_NS
```

```
+ if (de->parent == de)
+ de = &proc_root;
+#endif
i = filp->f_pos;
switch (i) {
case 0:
```
