
Subject: [PATCH 10/13] Make proc draw pids from appropriate namespace

Posted by [Pavel Emelianov](#) on Thu, 24 May 2007 13:04:10 GMT

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

The design is the following. The pids shown in proc tree are the ones get from the namespace the superblock belongs to.

As seen from the previous patch, when proc is mounted the current namespace is considered to be the owner of the superblock.

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

```
diff --git a/fs/proc/base.c b/fs/proc/base.c
index 8b426e9..e9811ec 100644
--- a/fs/proc/base.c
+++ b/fs/proc/base.c
@@ -2105,6 +2106,7 @@ struct dentry *proc_pid_lookup(struct in
     struct dentry *result = ERR_PTR(-ENOENT);
     struct task_struct *task;
     unsigned tgid;
+ struct pid_namespace *ns;

     result = proc_base_lookup(dir, dentry);
     if (!IS_ERR(result) || PTR_ERR(result) != -ENOENT)
@@ -2114,8 +2116,9 @@ struct dentry *proc_pid_lookup(struct in
     if (tgid == ~0U)
         goto out;

+ ns = (struct pid_namespace *)dentry->d_sb->s_fs_info;
     rcu_read_lock();
     task = find_task_by_pid(tgid);
+ task = find_task_by_pid_ns(tgid, ns);
     if (task)
         get_task_struct(task);
     rcu_read_unlock();
@@ -2132,7 +2135,8 @@ out:
     * Find the first task with tgid >= tgid
     *
     */
-static struct task_struct *next_tgid(unsigned int tgid)
+static struct task_struct *next_tgid(unsigned int tgid,
+ struct pid_namespace *ns)
{
    struct task_struct *task;
    struct pid *pid;
@@ -2140,9 +2144,9 @@ static struct task_struct *next_tgid(uns
```

```

    rcu_read_lock();
retry:
    task = NULL;
- pid = find_ge_pid(tgid);
+ pid = find_ge_pid(tgid, ns);
    if (pid) {
- tgid = pid->nr + 1;
+ tgid = pid_nr_ns(pid, ns) + 1;
        task = pid_task(pid, PIDTYPE_PID);
        /* What we to know is if the pid we have find is the
         * pid of a thread_group_leader. Testing for task
@@ -2182,6 +2186,7 @@ int proc_pid_readdir(struct file * filp,
        struct task_struct *reaper = get_proc_task(filp->f_path.dentry->d_inode);
        struct task_struct *task;
        int tgid;
+ struct pid_namespace *ns;

        if (!reaper)
            goto out_no_task;
@@ -2192,11 +2197,12 @@ int proc_pid_readdir(struct file * filp,
        goto out;
    }

+ ns = (struct pid_namespace *)filp->f_dentry->d_sb->s_fs_info;
    tgid = filp->f_pos - TGID_OFFSET;
- for (task = next_tgid(tgid);
+ for (task = next_tgid(tgid, ns);
        task;
-     put_task_struct(task), task = next_tgid(tgid + 1)) {
- tgid = task->pid;
+     put_task_struct(task), task = next_tgid(tgid + 1, ns)) {
+ tgid = __task_pid_nr_ns(task, ns);
        filp->f_pos = tgid + TGID_OFFSET;
        if (proc_pid_fill_cache(filp, dirent, filldir, task, tgid) < 0) {
            put_task_struct(task);
@@ -2324,6 +2330,7 @@ static struct dentry *proc_task_lookup(s
        struct task_struct *task;
        struct task_struct *leader = get_proc_task(dir);
        unsigned tid;
+ struct pid_namespace *ns;

        if (!leader)
            goto out_no_task;
@@ -2332,8 +2339,9 @@ static struct dentry *proc_task_lookup(s
        if (tid == ~0U)
            goto out;

+ ns = (struct pid_namespace *)dentry->d_sb->s_fs_info;

```

```

    rcu_read_lock();
- task = find_task_by_pid(tid);
+ task = find_task_by_pid_ns(tid, ns);
    if (task)
        get_task_struct(task);
    rcu_read_unlock();
@@ -2364,14 +2372,14 @@ out_no_task:
    * threads past it.
    */
    static struct task_struct *first_tid(struct task_struct *leader,
-    int tid, int nr)
+ int tid, int nr, struct pid_namespace *ns)
    {
        struct task_struct *pos;

        rcu_read_lock();
        /* Attempt to start with the pid of a thread */
        if (tid && (nr > 0)) {
- pos = find_task_by_pid(tid);
+ pos = find_task_by_pid_ns(tid, ns);
            if (pos && (pos->group_leader == leader))
                goto found;
        }
@@ -2440,6 +2448,7 @@ static int proc_task_readdir(struct file
    ino_t ino;
    int tid;
    unsigned long pos = filp->f_pos; /* avoiding "long long" filp->f_pos */
+ struct pid_namespace *ns;

    task = get_proc_task(inode);
    if (!task)
@@ -2473,12 +2482,13 @@ static int proc_task_readdir(struct file
    /* f_version caches the tgid value that the last readdir call couldn't
     * return. lseek aka telldir automatically resets f_version to 0.
     */
+ ns = (struct pid_namespace *)filp->f_dentry->d_sb->s_fs_info;
    tid = filp->f_version;
    filp->f_version = 0;
- for (task = first_tid(leader, tid, pos - 2);
+ for (task = first_tid(leader, tid, pos - 2, ns);
        task;
        task = next_tid(task), pos++) {
- tid = task->pid;
+ tid = __task_pid_nr_ns(task, ns);
    if (proc_task_fill_cache(filp, dirent, filldir, task, tid) < 0) {
        /* returning this tgid failed, save it as the first
         * pid for the next readir call */

```
