
Subject: Re: [PATCH 2/4] proc: Simplify proc_get_sb.
Posted by [Pavel Emelianov](#) on Wed, 12 Dec 2007 13:42:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

```
> The idle_thread now has a struct pid, so we can always find out know
> the pid of the child_reaper before we mount proc.
>
> Therefore we can remove the special cases for getting the pid of the
> child_reaper from proc_get_sb.
>
> Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>
> ---
> fs/proc/root.c | 17 +-----
> 1 files changed, 1 insertions(+), 16 deletions(-)
>
> diff --git a/fs/proc/root.c b/fs/proc/root.c
> index 81f99e6..f442967 100644
> --- a/fs/proc/root.c
> +++ b/fs/proc/root.c
> @@ -46,17 +46,6 @@ static int proc_get_sb(struct file_system_type *fs_type,
>  struct pid_namespace *ns;
>  struct proc_inode *ei;
>
> - if (proc_mnt) {
> - /* Seed the root directory with a pid so it doesn't need
> - * to be special in base.c. I would do this earlier but
> - * the only task alive when /proc is mounted the first time
> - * is the init_task and it doesn't have any pids.
> - */
> - ei = PROC_I(proc_mnt->mnt_sb->s_root->d_inode);
> - if (!ei->pid)
> - ei->pid = find_get_pid(1);
> - }
> -
> if (flags & MS_KERNMOUNT)
> ns = (struct pid_namespace *)data;
> else
> @@ -76,11 +65,7 @@ static int proc_get_sb(struct file_system_type *fs_type,
>  }
>
> ei = PROC_I(sb->s_root->d_inode);
> - if (!ei->pid) {
> - rcu_read_lock();
> - ei->pid = get_pid(find_pid_ns(1, ns));
> - rcu_read_unlock();
> - }
> + ei->pid = get_pid(ns->child_reaper);
```

That's not git-bisect safe - you move the child_reaper initialization before the call to prepare_proc only in the 4th patch.

```
> sb->s_flags |= MS_ACTIVE;  
> ns->proc_mnt = mnt;
```

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
