Subject: Re: [RFC][PATCH 14/16] Introduce proc_mnt for pid_ns Posted by Pavel Emelianov on Thu, 31 May 2007 11:48:21 GMT

View Forum Message <> Reply to Message

```
sukadev@us.ibm.com wrote:
> Subject: Introduce proc_mnt for pid_ns
> From: Dave Hansen <hansendc@us.ibm.com>
> The following patch completes the removal of the global proc mnt.
> It fetches the mnt on which to do dentry invalidations from the
> pid namespace in which the task appears.
>
> For now, there is only one pid namespace in mainline so this is
> straightforward. In the -lxc tree we'll have to do something
> more complex. The proc_flush_task() code takes a task, and
> needs to be able to find the corresponding proc superblocks on
> which that tasks's /proc/<pid> directories could appear. We
> can tell in which pid namespaces a task appears, so I put a
> pointer from the pid namespace to the corresponding proc mnt.
> /proc currently has some special code to make sure that the root
> directory gets set up correctly. It proc_mnt variable in order
> to find its way to the root inode.
>
> Signed-off-by: Dave Hansen <haveblue@us.ibm.com>
> Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
>
> fs/proc/base.c
                           > fs/proc/inode.c
                          11 ++++++
> fs/proc/root.c
                       > include/linux/pid_namespace.h | 1
> include/linux/proc_fs.h
> 5 files changed, 75 insertions(+), 22 deletions(-)
>
[snip]
> @ @ -2071,6 +2084,21 @ @ out:
> return;
> }
> +void proc_flush_task(struct task_struct *task)
> +{
> + int i;
> + struct pid *pid;
> + struct upid* upid;
```

```
> +
> + pid = task_pid(task);
> + if (!pid)
> + return;
The code below will never be called as task flushes all his pids
in __unhash_process() that happens before this. Or did I miss smth?
> + upid = &pid->upid_list[0];
> + for (i = 0; i < pid->num_upids; i++, upid++)
> + proc_flush_task_from_pid_ns(task, upid->pid_ns);
> +}
> +
> static struct dentry *proc_pid_instantiate(struct inode *dir,
        struct dentry * dentry,
        struct task_struct *task, const void *ptr)
>
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers
```