Subject: Re: [PATCH 15/15] Hooks over the code to show correct values to user
Posted by Oleg Nesterov on Tue, 31 Jul 2007 10:04:20 GMT
View Forum Message <> Reply to Message

On 07/30, Pavel Emelyanov wrote:
>
> Oleg Nesterov wrote:
> >On 07/26, Pavel Emelyanov wrote:
> >>int
> >>kill_proc(pid_t pid, int sig, int priv)
> >>{
> >>- return kill_proc_info(sig, __si_special(priv), pid);
> >>+ int ret;
> >>+
> >>+ rcu_read_lock();
> >>+ ret = kill_pid_info(sig, __si_special(priv), find_pid(pid));
> >>+ rcu_read_unlock();
> >>+ return ret;
> >>}
> >
> >I think this is wrong. kill_proc() should behave the same as
> >kill_proc_info(),
> >so this change is not needed. With this patch they use different namespaces
> >to find the task, this is not consistent.
>
> Actually, callers of this use tsk->pid (global pid) as an argument, so
> find_vpid() might return wrong value.

Yes I see. But still I don't agree on this issue.

kill_proc() is a simple wrapper on top of kill_proc_info(), not good
to break this. And with this patch they use different namespaces to
search the pid. Imho, not consistent.

Probably we can ignore this for now, but suppose we have some out-of-tree
driver which does kill_proc(pid_number), and the application from non-init
namespace does ioctl(SET_PID_NUMBER, getpid()).

And this is why btw I think find_pid/pid_nr should use active namespace, not
init_pid_ns. That driver can save "struct task_struct*" or "struct pid*".

OK, I understand it is a pain to "fix" the in-tree callers of kill_proc()
(say, we can introduce kill_pid_t() or something), so let's forget this.
In fact, we'd better remove kill_proc(), we should avoid using pid_t, the
callers should be converted to use struct pid.

Oleg.