Subject: Re: task_session() and task->signal->session
Posted by ebiederm on Fri, 17 Nov 2006 18:41:19 GMT
View Forum Message <> Reply to Message

Sukadev Bhattiprolu <sukadev@us.ibm.com> writes:

> Thanks.
>
> Eric W. Biederman [ebiederm@xmission.com] wrote:
> | Sukadev Bhattiprolu <sukadev@us.ibm.com> writes:
> |
> | > I am trying to understand the diff between task_session() and
> | > task->signal-session and why we have the two notions of session.
> | > Are we just phasing out one of them ?
> |
> | Yes.
>
> I guess we are phasing out task->signal->session.

Largely. There are a couple of cases where it makes sense to
optimize queries from the current pid namespace. Keeping
some of the pid_t values around for that case helps.

> |
> | > In general each process has a single session id. The only exception
> | > being a process that unshares its pid ns. Such process will have
> | > exactly two session ids, one for each pid ns. No ?
> |
> | Each process has a single session.
> |
> | That session can potentially have a different session id in each
> | pid namespace.
>
> both of us used the word "have" above. But do we actually store (in some
> data structure) the multiple session ids ? Or will the following work:
>
> Each task refers to another task (possibly itself) as its session
> leader (we find this using task_session()).
>
> The session leader, like any other task, has multiple process ids,
> one in each namespace.
>
> So to find the session id of a task, we find its session leader
> and find the appropriate process id of the session leader
>
> i.e we don't actually store the multiple sids a task

struct pid is that data structure. It just needs to be extended

a little to handle multiple pid namespaces.

Eric

_____