
Subject: Re: [PATCH 5/9] pid: Update pid_vnr to use task_active_pid_ns
Posted by [ebiederm](#) on Wed, 12 Dec 2007 14:20:19 GMT

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

> NAK. Your version will make more dereferences and one
> condition check (if (pid != NULL)). Besides, dead process
> can never call this thing - it is even never scheduled
> for execution.

Andrew please disregard this patch, it should not affect
the rest of the patch series.

Well EXIT_ZOMBIE is enough to cause us problems, not EXIT_DEAD,
so we can be scheduled.

Still we don't have any users currently and I actually don't expect us
to develop any, so this patch was just overkill right now.

I was thinking that there was a path to send_signal when
we were zombie with SEND_SIG_NOINFO, but that doesn't seem
to exist, and if there isn't one now we aren't likely
to grow one, and my signal patches remove the possibility
of that path being a problem.

Hmm. As to the cost argument.

Both pid->level and pid->numbers[pid->level].ns should be
on the same cache line so it should not be any real
extra memory reads. The branch should be predicted not
to happen so again it should not have a cost except a small
bit of code size.

So while I don't expect any real world cost this is the
common case so we don't want to put too many hoops in our
way either.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
