
Subject: Re: [PATCH] Mark timer_stats as incompatible with multiple pid namespaces

Posted by [ebiederm](#) on Thu, 13 Dec 2007 21:29:09 GMT

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

Ingo Molnar <mingo@elte.hu> writes:

> * Eric W. Biederman <ebiederm@xmission.com> wrote:

>

>> > the problem is, this interface stores historic PIDs too - i.e. PIDs

>> > of tasks that might have exited already.

>>

>> Well struct pid * works in that case if you grab the reference to it.

>

> but the display of the stats might happen much later. The point of this

> API is to save pid+comm, which gives users a good idea about what caused

> the events in the past - without having to pin any resource of that

> task.

Likewise struct pid is designed not to be a problem if pinned. It is a little heavier then it used to be with the addition of pid namespace support but not much. And if it is to heavy struct pid needs to be fixed.

Holding the struct pid very much does not pin the task struct, and it shouldn't pin any other resources. I agree 64bytes or so is a bit more to pin then 4 bytes but it really isn't a lot.

>> {

>> timer->start_site = NULL;

>> + pit_pid(timer->start_pid);

>> + timer->start_pid = NULL;

>

> s/pit/put, right?

Yes. Clearly I haven't tested it yet.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

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