

---

Subject: Re: [PATCH] Mark timer\_stats as incompatible with multiple pid namespaces

Posted by [Ingo Molnar](#) on Thu, 13 Dec 2007 13:04:23 GMT

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

---

\* Eric W. Biederman <[ebiederm@xmission.com](mailto:ebiederm@xmission.com)> wrote:

> > What the heck??? Please solve this properly instead of hiding it.  
> > /proc/timer\_stats is damn useful and it's a must-have for powertop  
> > to work.  
>  
> Hmm. Perhaps the dependency conflict should go in the other direction  
> then.  
>  
> My goal is to document the issue while a proper fix is being written.  
> I have known about this for all of about 1 day now. It was added  
> since last time I went through the kernel and made a thorough sweep of  
> pid users.  
>  
> What the proper fix is isn't even obvious at this point. Possibly it  
> is making /proc/timer\_stats disappear in child pid namespaces.  
> Which we don't currently have the infrastructure for.  
>  
> Possibly it is reworking the stats collection so we store a struct pid  
> \* instead of a pid\_t value. So we would know if the reader of the  
> value can even see processes you have collected stats for.

the problem is, this interface stores historic PIDs too - i.e. PIDs of tasks that might have exited already.

the proper way would be to also store the namespace address, and to filter out non-matching entries. (If leakage of this data across namespace creation is of any concern then flush out existing namespace data when a namespace is destroyed)

Ingo

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

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

---