
Subject: Re: pid namespace .text overhead

Posted by [Sukadev Bhattiprolu](#) on Tue, 11 Sep 2007 02:04:06 GMT

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

Cedric Le Goater [clg@fr.ibm.com] wrote:

| FYI,

|

| I just did a compile test on a 2.6.23-rc4-mm1 kernel with and without
| the following patches on a x86_64 defconfig (I also had to remove
| CONFIG_IPV6 for some compile reason) :

Thats a good point.

We have been a bit liberal with "inline" given that lot of the code was
going to affect the "fast-path" clone().

Should we shoot for a time/space trade-off or can we make some inlines
conditional (i.e inline when not configured for say embedded linux) ?

Suka

Containers mailing list

Containers@lists.linux-foundation.org

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

Subject: Re: pid namespace .text overhead

Posted by [Cedric Le Goater](#) on Tue, 11 Sep 2007 07:13:38 GMT

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

sukadev@us.ibm.com wrote:

> Cedric Le Goater [clg@fr.ibm.com] wrote:

> | FYI,

> |

> | I just did a compile test on a 2.6.23-rc4-mm1 kernel with and without
> | the following patches on a x86_64 defconfig (I also had to remove
> | CONFIG_IPV6 for some compile reason) :

>

> Thats a good point.

>

> We have been a bit liberal with "inline" given that lot of the code was
> going to affect the "fast-path" clone().

are are some figures :

with without

text data | text data | filename

1186	4	1099	4	kernel/capability.o
10390	0	9941	0	kernel/exit.o
10611	140	10434	140	kernel/fork.o
10765	72	10518	72	kernel/futex.o
957	56	939	56	kernel/nsproxy.o
3446	2292	2350	2228	kernel/pid.o
13930	45	13641	45	kernel/signal.o
10177	544	9819	544	kernel/sys.o
5083	56	4809	56	fs/proc/array.o
17097	184	16748	184	fs/proc/base.o
51471	1532	50474	1532	fs/proc/proc.o

> Should we shoot for a time/space trade-off or can we make some inlines
> conditional (i.e inline when not configured for say embedded linux) ?

I'm sure the embedded guys will appreciate :)

Thanks !

C.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: pid namespace .text overhead
Posted by [Dave Hansen](#) on Tue, 11 Sep 2007 15:36:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 2007-09-10 at 19:04 -0700, sukadev@us.ibm.com wrote:

>
> Thats a good point.
>
> We have been a bit liberal with "inline" given that lot of the code
> was
> going to affect the "fast-path" clone().
>
> Should we shoot for a time/space trade-off or can we make some inlines
> conditional (i.e inline when not configured for say embedded linux) ?

Can we make them switch somehow on -Os? That's what the embedded guys use.

-- Dave

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
