
Subject: Re: [PATCH 16/28] [FLAT 1/6] Changes in data structures for flat model
Posted by [Dave Hansen](#) on Tue, 19 Jun 2007 16:27:28 GMT

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

On Tue, 2007-06-19 at 11:18 +0400, Pavel Emelianov wrote:

> sukadev@us.ibm.com wrote:

> > Pavel Emelianov [xemul@openvz.org] wrote:

> > | This patch opens the flat model patches.

> > |

> > | The flat model idea is that struct pid has two numbers. The first one

> > | (pid->nr) is a global one and is unique in the system. The second one

> > | (pid->vnr) is a virtual pid. It is used on the kernel user boundary only.

> >

> > This approach duplicates 5 integers and 2 pointers per process for every

> > process in the system. While this may not be expensive for processes that

> > actually use multiple namespaces, doesn't it waste memory if majority of

> > processes exist only in one namespace ?

>

> task_struct alignment allows for it. so does the alignment of signal structure.

> and please note that this comes with appropriate ifdefs around. the only problem

> is with struct pid, but we're virtualizing it after all!

The #ifdefs do not help at all in the real world. Distributions will ship one and only one kernel configuration, save for things like SMP. This means that, no matter what, we're going to have the config option on.

So, for distribution customers, consider the #ifdef turned on all the time.

Containers mailing list

Containers@lists.linux-foundation.org

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