

---

Subject: Re: Re: [PATCH 16/28] [FLAT 1/6] Changes in data structures for flat model

Posted by [Pavel Emelianov](#) on Tue, 19 Jun 2007 07:52:29 GMT

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

---

sukadev@us.ibm.com wrote:

> Pavel Emelianov [xemul@openvz.org] 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!

>

> Hmm. I don't understand the last part "we are virtualizing 'struct pid'".

> Even so, with the FLAT model, every process will still have two

> pid\_t values, two hash-chain links etc - no ?

I mean that since we're adding some extra functionality to the kernel this is OK to extend the data structures. Moreover we have these changes under appropriate ifdefs.

> |

> | moreover - two integers and a pointer to the namespace is the minimal set of

> | fields for pid that is visible from two namespaces...

>

> I ignored the pid\_namespace pointer. But even a process that exists only

> in init\_pid\_ns would have the extra fields right ?

It will. But this is OK.

>

> \_\_\_\_\_  
> Containers mailing list

> Containers@lists.linux-foundation.org

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

>

---

Containers mailing list

