
Subject: Re: [RFC][PATCH 0/4] Object creation with a specified id
Posted by [Oren Laadan](#) on Fri, 14 Mar 2008 16:02:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Pavel Emelyanov wrote:

> Oren Laadan wrote:

>> Nadia Derbey wrote:

>>> Oren Laadan wrote:

>>>> Nadia.Derbey@bull.net wrote:

>>>>

>>>>> A couple of weeks ago, a discussion has started after Pierre's

>>>>> proposal for

>>>>> a new syscall to change an ipc id (see thread

>>>>> <http://lkml.org/lkml/2008/1/29/209>).

>>>>>

>>>>>

>>>>> Oren's suggestion was to force an object's id during its creation,

>>>>> rather

>>>>> than 1. create it, 2. change its id.

>>>>>

>>>>> So here is an implementation of what Oren has suggested.

>>>>>

>>>>> 2 new files are defined under /proc/self:

>>>>> . next_ipcid --> next id to use for ipc object creation

>>>>> . next_pids --> next upid nr(s) to use for next task to be forked

>>>>> (see patch #2 for more details).

>>>> Generally looks good. One meta-comment, though:

>>>>

>>>> I wonder why you use separate files for separate resources,

>>> That would be needed in a situation wheere we don't care about next,

>>> say, ipc id to be created but we need a predefined pid. But I must admit

>>> I don't see any pratical application to it.

>> exactly; why set the next-ipc value so far in advance ? I think it's

>> better (and less confusing) if we require that setting the next-id value

>> be done right before the respective syscall.

>

> And race with some other syscall caller? This will only work if the next-ipc-id

> and the next-pid are on a task_struct. Are they (at least supposed to be such)?

yes. that's the first detail I looked for in the patch :)

Containers mailing list

Containers@lists.linux-foundation.org

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