
Subject: Re: [RFC][PATCH 2.6.22-rc5] System V IPC: new IPC_SETID command to modify an ID

Posted by [Daniel Lezcano](#) on Wed, 20 Jun 2007 08:41:24 GMT

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

Kirill Korotaev wrote:

> Cedric Le Goater wrote:

>> Pierre Peiffer wrote:

>>

>>

>>> This patch adds a new IPC_SETID command to the System V IPCs set of
>>> commands, which allows to change the ID of an existing IPC.

>>>

>>> This command can be used through the semctl/shmctl/msgctl API, with the new
>>> ID passed as the third argument for msgctl and shmctl (instead of a
>>> pointer) and through the fourth argument for semctl.

>>>

>>> To be successful, the following rules must be respected:

>>> - the IPC exists

>>> - the user must be allowed to change the IPC attributes regarding the IPC
>>> permissions.

>>> - the new ID must satisfy the ID computation rule.

>>> - the entry (in the kernel internal table of IPCs) corresponding to the new
>>> ID must be free.

>>

>> That's an interesting way to reset the ids of sysv ipcs during a restart (after
>> a checkpoint) and we're looking for ways to do that among other things.

>>

>> How does it fit openvz ? Is it something openvz could use ?

>

> my personal imho is that we should not export such interfaces to user space
> and do the checkpointing from the kernel.

>

> it simplifies a lot of things and makes checkpointing more elegant.

> So until there is some user-space usage scenario of the patch I missed -

> i wouldn't commit it.

Hi Kirill,

I have no particular preference about checkpointing from userspace or kernelspace. Can you explain why doing the checkpoint from kernel simplifies the problem ? Do you have some examples about resources checkpointed from kernel which make the job easier ?

I thought we can divide system into subsystem and use the generic netlink to define a family for each subsystem. The result will be a checkpoint orchestrated from userspace but checkpointed from kernel.

-- Daniel

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