
Subject: Re: [PATCH 13/28] [PREP 13/14] Miscellaneous preparations in pid namespaces

Posted by [Pavel Emelianov](#) on Fri, 22 Jun 2007 01:10:04 GMT

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

sukadev@us.ibm.com wrote:

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

> | The most important one is moving exit_task_namespaces behind exit_notify

> | in do_exit() to make it possible to see the task's pid namespace to

> | properly notify the parent.

>

> Hmm. I think we tried this once a few months ago and got a crash in nfsd

> See <http://lkml.org/lkml/2007/1/17/75>

>

> [

> [

> [

> [

> [

> [

> [

> // exit_task_namespaces() after returning from exit_notify()

> [

> [

> [

>

> exit_notify() sets current->sighand to NULL and I think lockd_down() called

> from exit_task_namespaces/__put_mnt_ns() was accesssing current->sighand.

If sighand is set to NULL and then accessed then how is this related to pid namespace?

> Do your other patches in this set tweak something to prevent it ?

I think no. I'll check it for my current patches.

> Thats one of the reasons we had to remove pid_ns from nsproxy and use

> the pid_ns from pid->upid_list[i]->pid_ns.

>

> Suka

>

Containers mailing list

Containers@lists.linux-foundation.org

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