

---

Subject: Re: [PATCH 1/6] Define and use task\_active\_pid\_ns() wrapper  
Posted by [Sukadev Bhattiprolu](#) on Fri, 13 Jul 2007 06:03:17 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Pavel Emelianov [xemul@openvz.org] wrote:

| sukadev@us.ibm.com wrote:

| > Subject: [PATCH 1/6] Define and use task\_active\_pid\_ns() wrapper

| >

| > From: Sukadev Bhattiprolu <sukadev@us.ibm.com>

| >

| > With multiple pid namespaces, a process is known by some pid\_t in

| > every ancestor pid namespace. Every time the process forks, the

| > child process also gets a pid\_t in every ancestor pid namespace.

| >

| > While a process is visible in >=1 pid namespaces, it can see pid\_t's

| > in only one pid namespace. We call this pid namespace it's "active

| > pid namespace", and it is always the youngest pid namespace in which

| > the process is known.

| >

| > This patch defines and uses a wrapper to find the active pid namespace

| > of a process. The implementation of the wrapper will be changed in

| > when support for multiple pid namespaces are added.

| >

| > Changelog:

| > 2.6.22-rc4-mm2-pidns1:

| > - [Pavel Emelianov, Alexey Dobriyan] Back out the change to use

| > task\_active\_pid\_ns() in child\_reaper() since task->nsproxy

| > can be NULL during task exit (so child\_reaper() continues to

| > use init\_pid\_ns).

| >

| > to implement child\_reaper() since init\_pid\_ns.child\_reaper to

| > implement child\_reaper() since tsk->nsproxy can be NULL during exit.

| >

| > 2.6.21-rc6-mm1:

| > - Rename task\_pid\_ns() to task\_active\_pid\_ns() to reflect that a

| > process can have multiple pid namespaces.

| >

| > Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>

|

| Acked-by: Pavel Emelianov <xemul@openvz.org>

|

| But I'm not sure Andrew will like the construction like

|

| some\_function()->xxx;

|

| I used to do so in memory controller, but he asked to fix...

Hmm. I see this in the code:

```
rcu_dereference(p->real_parent)->tgid
```

Maybe rcu\_deref is a spl case. Well, if Andrew complains, I will add local vars as needed. I think its easier to read with a wrapper than expanded out.

Thanks,

Suka

---

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

---