

---

Subject: Re: [PATCH 2/3] Pid ns helpers for signals  
Posted by [Sukadev Bhattiprolu](#) on Mon, 03 Sep 2007 16:01:47 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Oleg Nesterov [oleg@tv-sign.ru] wrote:

| On 09/01, Oleg Nesterov wrote:

| >

| > On 08/31, sukadev@us.ibm.com wrote:

| > >

| > > +static struct pid\_namespace \*get\_task\_pid\_ns(struct task\_struct \*tsk)

| > > +{

| > > + struct pid \*pid;

| > > + struct pid\_namespace \*ns;

| > > +

| > > + pid = get\_task\_pid(tsk, PIDTYPE\_PID);

| > > + ns = get\_pid\_ns(pid\_active\_ns(pid));

| > > + put\_pid(pid);

| > > +

| > > + return ns;

| > > +}

| >

| > Hmm. Firstly, we don't need this for the "current", but all users of this func

| > also do get\_task\_pid\_ns(current).

| >

| > Also, we don't need get/put\_pid. rcu locks are enough,

| >

| > rcu\_read\_lock();

| > ns = get\_pid\_ns(pid\_active\_ns(task\_pid(tsk)));

| > rcu\_read\_unlock();

| >

| > However, do we really need this complications right now? Currently, we use

| > this "compare namespaces" helpers only when we know that "struct pid" is

| > stable. We are sending the signal to that task, it must be pid\_alive(), and

| > we either locked the task itself, or we hold tasklist.

|

| (forgot to mention)

|

| Otherwise, it is not safe to use "tsk" in get\_task\_pid\_ns(), so I don't think

| these get/put pid/pid\_ns games make too much sense.

get\_pid(), put\_pid(), get\_pid\_ns(), put\_pid\_ns() all allow pid to be NULL.

You mean tsk itself can be NULL bc task is exiting ?

|

| Oleg.

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

