
Subject: Re: [RFC][PATCH 14/16] Introduce proc_mnt for pid_ns
Posted by [Sukadev Bhattiprolu](#) on Thu, 31 May 2007 19:50:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dave Hansen [hansendc@us.ibm.com] wrote:

| On Thu, 2007-05-31 at 15:48 +0400, Pavel Emelianov wrote:

| > > +void proc_flush_task(struct task_struct *task)

| > > +{

| > > + int i;

| > > + struct pid *pid;

| > > + struct upid* upid;

| > > +

| > > + pid = task_pid(task);

| > > + if (!pid)

| > > + return;

| >

| > The code below will never be called as task flushes all his pids

| > in __unhash_process() that happens before this. Or did I miss smth?

|

| Nope, that's a very nice catch. Suka is working on a patch to fix this

| right now. We just need to keep the list of pids accessible a little

| longer, but I think Suka ran into a chicken-and-egg problem while trying

| to solve this. Suka, care to post your workaround?

Yes. the chicken and egg problem is in the order of the __unhash_process()
and proc_flush_task().

With __unhash_process() first, proc_flush_task() has no way of knowing
the pid namespaces the process belongs to. With proc_flush_task() first,
we find the process but have its dentry removed from proc. This crashes
when I run a test in a tight loop.

A quick/dirty fix is to save the pid namespace list before __unhash_process()
and use that in proc_flush_task(). That would make the exit() path quite
expensive. I am still investigating the crash and looking for a better option.

|
| -- Dave

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