Subject: Re: [PATCH][usercr]: Ghost tasks must be detached
Posted by Sukadev Bhattiprolu on Thu, 10 Feb 2011 17:54:09 GMT
View Forum Message <> Reply to Message

Louis Rilling [Louis.Rilling@kerlabs.com] wrote:
| > I can reproduce a crash with 2.6.32 - where if container-init terminates
| > before a detached child, we get a crash when the detached child calls
| > proc_flush_mnt(). I suspected it was because do_wait_thread() skipped
| > over detached tasks (in 2.6.32).
| >
| > The same test case does not crash on 2.6.37 - which includes the above commit.
| > The removes the check for detached tasks, my initial guess is that the above
| > commit, may have contributed to _fixing_ the crash in 2.6.37.
|
| Hm, I don't see how this commit changed things for detached tasks, unless ptrace
| is involved. Detached tasks go atomically
| from ->exit_state == 0 to ->exit_state == EXIT_DEAD in exit_notify(),
| because tracehook_notify_death() returns DEATH_REAP for all not ptraced detached
| tasks.
|
| What do you think has changed precisely?

Well, one of the changes in the commit is this:

```
@@ -1551,14 +1554,9 @@ static int do_wait_thread(struct wait_opts *wo, struct task_struct *tsk)
        struct task_struct *p;

        list_for_each_entry(p, &tsk->children, sibling) {
-               /*
-                * Do not consider detached threads.
-                */
-               if (!task_detached(p)) {
-                       int ret = wait_consider_task(wo, 0, p);
-                       if (ret)
-                               return ret;
-               }
+               int ret = wait_consider_task(wo, 0, p);
+               if (ret)
+                       return ret;
        }

        return 0;
```

If it was a detached task, do_wait_thread() skipped over it. In the C/R
kernel we were setting the ->exit_signal to -1 for a "ghost" process.
I assumed that the container-init exited without waiting for the "ghost"
and we were getting the crash in proc_flush_mnt() when the ghost exited.

Sukadev

_____

Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containe rs