
Subject: [PATCH 11/20] Move alloc_pid() lower in copy_process()
Posted by [Pavel Emelianov](#) on Tue, 07 Aug 2007 09:29:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

When we create new namespace we will need to allocate the struct pid, that will have one extra struct upid in array, comparing to the parent.

Thus we need to know the new namespace (if any) in alloc_pid() to init this struct upid properly, so move the alloc_pid() call lower in copy_process().

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

fork.c | 31 ++++++-----
1 files changed, 16 insertions(+), 15 deletions(-)

```
--- ./kernel/fork.c.ve12 2007-08-06 12:44:54.000000000 +0400
+++ ./kernel/fork.c 2007-08-06 12:44:57.000000000 +0400
@@ -1028,16 +1027,9 @@ static struct task_struct *copy_process(
     if (p->binfmt && !try_module_get(p->binfmt->module))
         goto bad_fork_cleanup_put_domain;

- if (pid != &init_struct_pid) {
-     pid = alloc_pid(task_active_pid_ns(p));
-     if (!pid)
-         goto bad_fork_put_binfmt_module;
- }
-
    p->did_exec = 0;
    delayacct_tsk_init(p); /* Must remain after dup_task_struct() */
    copy_flags(clone_flags, p);
- p->pid = pid_nr(pid);
    INIT_LIST_HEAD(&p->children);
    INIT_LIST_HEAD(&p->sibling);
    p->vfork_done = NULL;
@@ -1112,10 +1104,6 @@ static struct task_struct *copy_process(
    p->blocked_on = NULL; /* not blocked yet */
    #endif

- p->tgid = p->pid;
- if (clone_flags & CLONE_THREAD)
-     p->tgid = current->tgid;
-
    if ((retval = security_task_alloc(p))
        goto bad_fork_cleanup_policy;
    if ((retval = audit_alloc(p))
```

```

@@ -1141,6 +1129,18 @@ static struct task_struct *copy_process(
    if (retval)
        goto bad_fork_cleanup_namespaces;

+ if (pid != &init_struct_pid) {
+     retval = -ENOMEM;
+     pid = alloc_pid(task_active_pid_ns(p));
+     if (!pid)
+         goto bad_fork_cleanup_namespaces;
+ }
+
+ p->pid = pid_nr(pid);
+ p->tgid = p->pid;
+ if (clone_flags & CLONE_THREAD)
+     p->tgid = current->tgid;
+
+ p->set_child_tid = (clone_flags & CLONE_CHILD_SETTID) ? child_tidptr : NULL;
+ /*
+  * Clear TID on mm_release()?
@@ -1237,7 +1237,7 @@ static struct task_struct *copy_process(
    spin_unlock(&current->sigband->siglock);
    write_unlock_irq(&tasklist_lock);
    retval = -ERESTARTNOINTR;
- goto bad_fork_cleanup_namespaces;
+ goto bad_fork_free_pid;
}

if (clone_flags & CLONE_THREAD) {
@@ -1283,6 +1294,9 @@ static struct task_struct *copy_process(
    container_post_fork(p);
    return p;

+bad_fork_free_pid:
+ if (pid != &init_struct_pid)
+     free_pid(pid);
bad_fork_cleanup_namespaces:
    exit_task_namespaces(p);
bad_fork_cleanup_keys:
@@ -1322,9 +1329,6 @@ bad_fork_cleanup_container:
#endif
    container_exit(p, container_callbacks_done);
    delayacct_tsk_free(p);
- if (pid != &init_struct_pid)
-     free_pid(pid);
-bad_fork_put_binfmt_module:
    if (p->binfmt)
        module_put(p->binfmt->module);
bad_fork_cleanup_put_domain:

```
