
Subject: Re: [PATCH 0/3] Make tasks always have non-zero pids
Posted by [Sukadev Bhattiprolu](#) on Fri, 05 Oct 2007 05:53:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

Pavel Emelianov [xemul@openvz.org] wrote:
| Some time ago Sukadev noticed that the vmlinux size has

Cedric pointed it out to me first :-)

| grown 5Kb due to merged pid namespaces. One of the big
| problems with it was fat inline functions. The other thing
| was noticed by Matt - the checks for task's pid to be not
| NULL take place and make the kernel grow due to inlining,
| but these checks are not always needed.

| In this series I introduce a static pid (dummy), according
| to Matt's proposal, which is assigned to tasks during the
| detach_pid and transfer_pid instead of NULL. This pid lives
| in the init pid namespace and has the id = 0, so all the
| task_xid_xnr() calls will still return 0 on a dead task.

| Places that get the struct pid from task either get it from
| the current (in this case they will never get this dummy),
| or use it to compare with some other value (so they will
| work the same for both NULL and dummy pids).

| This saves up to 340 bytes for i386 kernel with minimal
| config and probably more with more users of pids.

| Tested on i386 and x86_64 boxes. Tasks still live and die,
| namespaces and proc still work.

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

Acked-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
