

---

Subject: Re: [PATCH 0/16] Pid namespaces

Posted by [Pavel Emelianov](#) on Mon, 16 Jul 2007 08:47:10 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

> My x86\_64 system boots fine but crashes as below, when I run my  
> 'pidns\_exec' test with a simple program that prints getpid(), getppid()  
> etc of the process in the child pid ns.  
>  
> Pls see  
>  
> <http://www.geocities.com/sukadevb/Pidspace/2.6.22-rc6-mm1-pavel-1.tgz>  
>  
> for the patches I currently have applied and let me know if I need more  
> on top.  
>  
> And see  
>  
> <http://www.geocities.com/sukadevb/Pidspace/Test1/>  
>  
> for the test programs. You may need to run the 'mypid-loop.x' script  
> to repro the crash. The pidns\_exec.c program calls clone() with CLONE\_NEWPID  
> and execs the given program (it was included in Patch 0 of the patchset I  
> posted to Containers).  
>  
> Suka  
>  
> login: Unable to handle kernel NULL pointer dereference at 00000000000002fc RIP:  
> [[ffffffffff802b9e5e](#)] proc\_get\_sb+0xfb/0x138  
> PGD 104d53067 PUD 104d4d067 PMD 0  
> Oops: 0002 [1] SMP  
> CPU 2  
> Modules linked in:  
> Pid: 3279, comm: pidns\_exec Not tainted 2.6.22-rc6-mm1-ovz1 #10  
> RIP: 0010:[[ffffffffff802b9e5e](#)] [[ffffffffff802b9e5e](#)] proc\_get\_sb+0xfb/0x138  
> RSP: 0018:ffff8101029d7d28 EFLAGS: 00010202  
> RAX: ffff810100651840 RBX: ffff810104461400 RCX: ffff810100651878  
> RDX: 0000000000000000 RSI: ffffffff806e5460 RDI: 0000000000000238  
> RBP: ffff810102d886f8 R08: ffff810104461400 R09: ffff810100026000  
> R10: 0000000000000000 R11: 0000000000000002 R12: ffff8101029c6000  
> R13: 0000000002000000 R14: ffffffff806ee920 R15: ffff810102088cc0  
> FS: 00002b0b499ec6f0(0000) GS:ffff81010069c3c0(0000) knlGS:0000000000000000  
> CS: 0010 DS: 0000 ES: 0000 CR0: 000000008005003b  
> CR2: 00000000000002fc CR3: 000000010381b000 CR4: 000000000000006e0  
> DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000  
> DR3: 0000000000000000 DR6: 00000000ffff0ff0 DR7: 0000000000000400  
> Process pidns\_exec (pid: 3279, threadinfo ffff8101029d6000, task ffff81010269e7f0)  
> Stack: ffff810102088cc0 ffff810102088cc0 00000000ffffff4 ffffffff806ee920  
> ffffffff8065f9d9 ffff8101029c6000 0000000002000000 ffffffff80287164

```

> 00000000000000d0 ffff8101029c6000 ffffffff806e5460 ffff8101029c6000
> Call Trace:
> [<ffffffff80287164>] vfs_kern_mount+0x4f/0x8b
> [<ffffffff802b9cf4>] pid_ns_prepare_proc+0x13/0x2e
> [<ffffffff80245be3>] copy_pid_ns+0xd7/0x164
> [<ffffffff8024af34>] create_new_namespaces+0xde/0x192
> [<ffffffff8024b0aa>] copy_namespaces+0x4b/0x85
> [<ffffffff802347e2>] copy_process+0xcb4/0x1439
> [<ffffffff8020bbee>] system_call+0x7e/0x83
> [<ffffffff8023556a>] do_fork+0x6c/0x1e7
> [<ffffffff8020bf07>] ptregscall_common+0x67/0xb0
>

```

Here's the patch fixing the problem. So, Suka, I propose that you review my patches, point out things that you don't like and would like to see your code instead. After all I will re-split the set with your fixes, mark some patches with your From: and send them to Andrew. What do you think?

---

```

--- ./fs/proc/root.c.procpidnsfix 2007-07-16 10:32:00.000000000 +0400
+++ ./fs/proc/root.c 2007-07-16 12:34:35.000000000 +0400
@@ -79,8 +79,6 @@ static int proc_get_sb(struct file_syste
    if (!ei->pid)
        ei->pid = find_get_pid(1);
    sb->s_flags |= MS_ACTIVE;
-
- mntput(ns->proc_mnt);
- ns->proc_mnt = mnt;
}

```

---

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

---