
Subject: Re: nptl perf bench and profiling with pidns patchsets
Posted by [Sukadev Bhattiprolu](#) on Sat, 09 Jun 2007 18:47:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

Alexey Dobriyan [adobriyan@sw.ru] wrote:

| On Sat, Jun 09, 2007 at 12:10:25PM +0400, Pavel Emelianov wrote:
| > > * definitely better results for suka's patchset. suka's patchset is
| > > also getting better results with unixbench on a 2.6.22-rc1-mm1 but
| > > the values are really dispersed. can you confirm ?
| > > * suka's patchset would benefit from some optimization in init_upid()
| > > and dup_struct_pid()
| >
| > We have found the reason why Suka's patches showed better performance.
| > Some time ago I sent a letter saying that proc_flush_task() actually
| > never worked with his patches - that's the main problem. After removing
| > this call from my patches the results turned to those similar to my.

i.e with the call removed from both our sets, my patchset is about 1-1.5%
slower than yours ?

| >
| > I'd also like to note that broken-out set of patches is not git bisect
| > safe at all. The very first patch of his own OOPSes the node.
|
| FWIW, it's EIP is at forget_original_parent+0x25 on boot
| Process: khelper
| exit_notify
| do_exit
| copy_vm86_regs_to_user
| kernel_execve
| ____call_usermodehelper

Thanks for pointing it out. I will backout this change from patch #1 bc
tsk->nsproxy can be null during exit.

```
static inline struct task_struct *child_reaper(struct task_struct *tsk)
{
-   return init_pid_ns.child_reaper;
+   return task_active_pid_ns(tsk)->child_reaper;
}
```

I also fixed the problem in proc_flush_task() and am working on fixing
signals. After that I will port to more recent kernel and ensure they
are bisect safe.

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