
Subject: pid namespace .text overhead

Posted by [Cedric Le Goater](#) on Mon, 10 Sep 2007 19:38:19 GMT

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

FYI,

I just did a compile test on a 2.6.23-rc4-mm1 kernel with and without the following patches on a x86_64 defconfig (I also had to remove CONFIG_IPV6 for some compile reason) :

```
+ pid-namespaces-rework-forget_original_parent.patch
+ pid-namespaces-move-exit_task_namespaces.patch
+ pid-namespaces-introduce-ms_kernmount-flag.patch
+ pid-namespaces-prepare-proc_flust_task-to-flush-entries-from-multiple-proc-trees.patch
+ pid-namespaces-introduce-struct-upid.patch
+ pid-namespaces-add-support-for-pid-namespaces-hierarchy.patch
+ pid-namespaces-make-alloc_pid-free_pid-and-put_pid-work-with-struct-upid.patch
+ pid-namespaces-helpers-to-obtain-pid-numbers.patch
+ pid-namespaces-helpers-to-find-the-task-by-its-numerical-ids.patch
+ pid-namespaces-helpers-to-find-the-task-by-its-numerical-ids-fix.patch
+ pid-namespaces-move-alloc_pid-lower-in-copy_process.patch
+ pid-namespaces-make-proc-have-multiple-superblocks-one-for-each-namespace.patch
+ pid-namespaces-miscellaneous-preparations-for-pid-namespaces.patch
+ pid-namespaces-allow-cloning-of-new-namespace.patch
+
pid-namespaces-allow-cloning-of-new-namespace-fix-check-for-return-value-of-create_pid_name
space.patch
+ pid-namespaces-make-proc_flush_task-actually-from-entries-from-multiple-namespaces.patch
+ pid-namespaces-initialize-the-namespaces-proc_mnt.patch
+ pid-namespaces-create-a-slab-cache-for-struct-pid_namespace.patch
+ pid-namespaces-allow-signalling-container-init.patch
+ pid-namespaces-destroy-pid-namespace-on-init-death.patch
+ pid-namespaces-changes-to-show-virtual-ids-to-user.patch
+
pid-namespaces-changes-to-show-virtual-ids-to-user-fix-the-return-value-of-sys_set_tid_address.
patch
+
pid-namespaces-changes-to-show-virtual-ids-to-user-use-find_task_by_pid_ns-in-places-that-oper
ate-with-virtual.patch
+
pid-namespaces-changes-to-show-virtual-ids-to-user-use-find_task_by_pid_ns-in-places-that-oper
ate-with-virtual-fix.patch
+
pid-namespaces-changes-to-show-virtual-ids-to-user-use-find_task_by_pid_ns-in-places-that-oper
ate-with-virtual-fix-2.patch
+
pid-namespaces-changes-to-show-virtual-ids-to-user-use-find_task_by_pid_ns-in-places-that-oper
ate-with-virtual-fix-3.patch
```

```

+
pid-namespaces-changes-to-show-virtual-ids-to-user-sys_getsid-sys_getpgid-return-wrong-id-for-t
ask-from-another.patch
+
pid-namespaces-changes-to-show-virtual-ids-to-user-fix-the-sys_setpgrp-to-work-between-names
paces.patch
+ pid-namespaces-changes-to-show-virtual-ids-to-user-fix.patch
+ pid-namespaces-remove-the-struct-pid-unneeded-fields.patch
+ isolate-some-explicit-usage-of-task-tgid.patch
+ isolate-some-explicit-usage-of-task-tgid-fix.patch
+ isolate-some-explicit-usage-of-task-tgid-fix-fix.patch

```

I got a less than *6k* difference in .text. Here are the detailed size results.

size without:

text	data	bss	dec	hex	filename
5067718	892082	725544	6685344	6602a0	vmlinux

size with:

text	data	bss	dec	hex	filename
5073314	892210	725544	6691068	6618fc	vmlinux

section	without	with
.text	3383305	3388569
__ex_table	16704	16704
__bug_table	41136	41208
.rodata	1276129	1276469
.pci_fixup	4032	4032
__ksymtab	43984	43984
__ksymtab_gpl	13024	13056
__ksymtab_gpl_future	48	48
__ksymtab_strings	83388	83420
__param	7800	7800
.data	482928	483056
.data.cacheline_aligned	187008	187008
.data.read_mostly	41600	41600
.vsyscall_0	227	227
.vsyscall_fn	54	54
.vsyscall_gtod_data	80	80
.vsyscall_1	52	52
.vsyscall_2	91	91
.vgetcpu_mode	4	4
.jiffies	8	8
.vsyscall_3	8	8
.data.init_task	8192	8192
.data.page_aligned	4096	4096

.smp_locks	33584	33624
.init.text	153713	153529
.init.data	137714	137714
.init.setup	3168	3168
.initcall.init	2200	2200
.con_initcall.init	16	16
.altinstructions	1243	1243
.altinstr_replacement	244	244
.exit.text	4739	4739
.note	24	24
.vds0	4056	4056
.init.ramfs	133	133
.data_percpu	25064	25064
.data_nosave	4	4
.bss	725544	725544
.comment	44352	44352
Total	6729696	6735420

kernel/pid.o .text increases a lot (more than a 1K) but some inline would also need to be checked (that i didn't do yet)

Cheers,

C.

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