
Subject: Re: Kernel text size with the pid namespaces fixes
Posted by [Pavel Emelianov](#) on Tue, 25 Sep 2007 14:28:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Cedric Le Goater wrote:

> Hello Pavel !

>

> Pavel Emelyanov wrote:

>> Hi, Suka

>>

>> I have posted tree patches with un-inlining of some macros

>> and static inline calls that are expanded into another call.

>>

>> At my i386 notebook this saved about half-a-kilo from the

>> vmlinux. Could you please check these patches at your

>> boxes/configs to see how much progress we have here.

>>

>> Here's the combined patch for 2.6.24-rc8-mm1

>

> on my qemu configuration, scripts/bloat-o-meter gave the following

> results.

>

> it would be interesting to run some micro benchmark. Did you have

> time to do that ?

We did. We ran an nptl test and it showed no performance degradation.

This was expected - I didn't introduce new "call"-s in this patch. See,

all the uninlined macros called some function at the end. I just moved

the whole code (with the preparations) into a call.

> Thanks,

>

> C.

>

> add/remove: 9/0 grow/shrink: 3/47 up/down: 194/-683 (-489)

> function old new delta

> find_task_by_vpid - 29 +29

> find_vpid - 25 +25

> task_tgid_nr_ns - 22 +22

> task_session_nr_ns - 22 +22

> task_pgrp_nr_ns - 22 +22

> find_task_by_pid - 19 +19

> task_pid_nr_ns - 16 +16

> find_task_by_pid_ns - 16 +16

> find_pid - 15 +15

> proc_task_readdir 762 766 +4

> msg_exit_ns 117 119 +2

> _request_firmware 688 690 +2

> nfs_readpage_truncate_uninitialised_page	298	297	-1
> sys_swapoff	1898	1896	-2
> set_get_cmap	332	330	-2
> __iovec_copy_from_user_inatomic	88	86	-2
> proc_pid_readdir	438	434	-4
> init_pit_clocksource	82	78	-4
> do_notify_parent_cldstop	297	293	-4
> do_notify_parent	394	390	-4
> kill_proc	39	34	-5
> sys_timer_create	720	713	-7
> do_setlink	796	789	-7
> debug_rt_mutex_print_deadlock	406	399	-7
> check_clock	122	115	-7
> rest_init	78	69	-9
> posix_cpu_clock_get	221	212	-9
> __set_special_pids	154	144	-10
> vma_adjust	855	844	-11
> tty_ioctl	3274	3263	-11
> kthreadd	248	237	-11
> sys_kill	313	301	-12
> sys_capget	285	273	-12
> sys_getsid	128	114	-14
> sys_getpgid	128	114	-14
> posix_cpu_timer_create	240	226	-14
> kill_proc_info	45	31	-14
> find_get_pid	31	16	-15
> sys_sched_getscheduler	92	76	-16
> sys_sched_getparam	143	127	-16
> sys_get_robust_list	149	133	-16
> sched_setaffinity	259	243	-16
> sched_getaffinity	114	98	-16
> ptrace_get_task_struct	90	74	-16
> futex_lock_pi	2435	2419	-16
> f_setown	68	52	-16
> do_tkill	280	264	-16
> do_sched_setscheduler	113	97	-16
> sys_sched_rr_get_interval	161	141	-20
> sys_capset	702	681	-21
> eligible_child	215	191	-24
> do_task_stat	2135	2107	-28
> sys_setpriority	446	416	-30
> sys_setpgid	419	389	-30
> sys_ioprio_set	471	441	-30
> sys_ioprio_get	388	358	-30
> sys_getpriority	439	409	-30
> proc_pid_status	1216	1186	-30
> do_wait	2677	2639	-38
>			

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
