Subject: [PATCH] cgroup: anotate two variables with __read_mostly
Posted by Li Zefan on Thu, 05 Jun 2008 06:02:47 GMT
View Forum Message <> Reply to Message

- need_forkexit_callback will be read only after system boot.
- use_task_css_set_links will be read only after it's set.

And these 2 variables are checked when a new process is forked.

Signed-off-by: Li Zefan <lizf@cn.fujitsu.com>
---
 kernel/cgroup.c | 4 ++--
 1 files changed, 2 insertions(+), 2 deletions(-)

diff --git a/kernel/cgroup.c b/kernel/cgroup.c
index 15ac0e1..f5a1273 100644
--- a/kernel/cgroup.c
+++ b/kernel/cgroup.c
@@ -118,7 +118,7 @@ static int root_count;
  * extra work in the fork/exit path if none of the subsystems need to
  * be called.
  */
-static int need_forkexit_callback;
+static int need_forkexit_callback __read_mostly;
 static int need_mm_owner_callback __read_mostly;

 /* convenient tests for these bits */
@@ -220,7 +220,7 @@ static struct hlist_head *css_set_hash(struct cgroup_subsys_state *css[])
  * task until after the first call to cgroup_iter_start(). This
  * reduces the fork()/exit() overhead for people who have cgroups
  * compiled into their kernel but not actually in use */
-static int use_task_css_set_links;
+static int use_task_css_set_links __read_mostly;

 /* When we create or destroy a css_set, the operation simply
  * takes/releases a reference count on all the cgroups referenced
--
1.5.4.rc3

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