
Subject: [PATCH] sched: don't call task_group() many times in set_task_rq()

Posted by [Andrew Vagin](#) on Fri, 11 Nov 2011 14:04:00 GMT

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

It improves performance, especially if autogroup enabled.

The size of set_task_rq() was 0x180 and now it is 0xa0.

Signed-off-by: Andrew Vagin <avagin@openvz.org>

kernel/sched.c | 12 ++++++-----

1 files changed, 8 insertions(+), 4 deletions(-)

diff --git a/kernel/sched.c b/kernel/sched.c

index d87c6e5..a740eec 100644

--- a/kernel/sched.c

+++ b/kernel/sched.c

@@ -766,14 +766,18 @@ static inline struct task_group *task_group(struct task_struct *p)

/* Change a task's cfs_rq and parent entity if it moves across CPUs/groups */

static inline void set_task_rq(struct task_struct *p, unsigned int cpu)

{

+ #if defined(CONFIG_FAIR_GROUP_SCHED) || defined(CONFIG_RT_GROUP_SCHED)

+ struct task_group *tg = task_group(p);

+ #endif

+

 #ifdef CONFIG_FAIR_GROUP_SCHED

- p->se.cfs_rq = task_group(p)->cfs_rq[cpu];

- p->se.parent = task_group(p)->se[cpu];

+ p->se.cfs_rq = tg->cfs_rq[cpu];

+ p->se.parent = tg->se[cpu];

 #endif

 #ifdef CONFIG_RT_GROUP_SCHED

- p->rt.rt_rq = task_group(p)->rt_rq[cpu];

- p->rt.parent = task_group(p)->rt_se[cpu];

+ p->rt.rt_rq = tg->rt_rq[cpu];

+ p->rt.parent = tg->rt_se[cpu];

 #endif

}

--

1.7.1

Subject: Re: [PATCH] sched: don't call task_group() many times in set_task_rq()

Posted by [Peter Zijlstra](#) on Sat, 12 Nov 2011 09:21:58 GMT

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

On Fri, 2011-11-11 at 13:26 -0800, Paul Turner wrote:
> Signed-off-by: Paul Turner <pjt@google.com>

I'll take that as Acked-by, ok?
