

---

Subject: Re: linux-next: lockdep whinge in cgroup\_rmdir  
Posted by [Nick Piggin](#) on Fri, 14 Jan 2011 04:10:31 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Fri, Jan 14, 2011 at 3:06 PM, Li Zefan <lizf@cn.fujitsu.com> wrote:

> Nick Piggin wrote:

>> It looks like it is just a missing parent->child lock order annotation, but

>> mainline cgroupfs code looks to be OK there. What does

>> cgroup\_clear\_directory() look like in mmotm?

>

> It's not from cgroup\_clear\_directory()..

>

> This should fix it:

Yep, thanks!

> =====

>

> From: Li Zefan <lizf@cn.fujitsu.com>

> Date: Fri, 14 Jan 2011 11:34:34 +0800

> Subject: [PATCH] cgroups: Fix a lockdep warning at cgroup removal

>

> Commit 2fd6b7f5 ("fs: dcache scale subdirs") forgot to annotate a dentry

> lock, which caused a lockdep warning.

>

> Reported-by: Valdis Kletnieks <Valdis.Kletnieks@vt.edu>

> Signed-off-by: Li Zefan <lizf@cn.fujitsu.com>

> ---

> kernel/cgroup.c | 2 +-  
> 1 files changed, 1 insertions(+), 1 deletions(-)

>

> diff --git a/kernel/cgroup.c b/kernel/cgroup.c  
> index 5c5f4cc..db983e2 100644

> --- a/kernel/cgroup.c

> +++ b/kernel/cgroup.c

> @@ -910,7 +910,7 @@ static void cgroup\_d\_remove\_dir(struct dentry \*dentry)

>

> parent = dentry->d\_parent;

> spin\_lock(&parent->d\_lock);

> - spin\_lock(&dentry->d\_lock);

> + spin\_lock\_nested(&dentry->d\_lock, DENTRY\_D\_LOCK\_NESTED);

> list\_del\_init(&dentry->d\_u.d\_child);

> spin\_unlock(&dentry->d\_lock);

> spin\_unlock(&parent->d\_lock);

> --

> 1.7.3.1

>

>

---

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

---