
Subject: [PATCH v2 2/5] blkcg: protect blkcg->policy_list
Posted by [Glauber Costa](#) on Mon, 23 Apr 2012 19:37:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

policy_list walks are protected with blkcg->lock everywhere else in the code. In destroy(), they are not. Because destroy is usually protected with the cgroup_mutex(), this is usually not a problem. But it would be a lot better not to assume this.

Signed-off-by: Glauber Costa <glommer@parallels.com>

block/blk-cgroup.c | 2 ++
1 files changed, 2 insertions(+), 0 deletions(-)

```
diff --git a/block/blk-cgroup.c b/block/blk-cgroup.c
index 126c341..35fd701 100644
--- a/block/blk-cgroup.c
+++ b/block/blk-cgroup.c
@@ -1557,10 +1557,12 @@ static void blkicg_destroy(struct cgroup *cgroup)
     spin_unlock(&blkio_list_lock);
 } while (1);

+ spin_lock_irqsave(&blkcg->lock, flags);
 list_for_each_entry_safe(pn, pntmp, &blkcg->policy_list, node) {
     blkio_policy_delete_node(pn);
     kfree(pn);
 }
+ spin_unlock_irqrestore(&blkcg->lock, flags);

 free_css_id(&blkio_subsys, &blkcg->css);
 rcu_read_unlock();
--
1.7.7.6
```
