
Subject: [RFC][for -mm] memory cgroup enhancements take3 [7/9] add pre_destroy

Posted by [KAMEZAWA Hiroyuki](#) on Tue, 30 Oct 2007 11:20:52 GMT

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

My main purpose of this patch is for memory controller..

This patch adds a handler "pre_destroy" to cgroup_subsys.
It is called before cgroup_rmdir() checks all subsys's refcnt.

I think this is useful for subsyses which have some extra refs even if there are no tasks in cgroup. By adding pre_destroy(), the kernel keeps the rule "destroy() against subsystem is called only when refcnt=0." and allows css's ref to be used by other objects than tasks.

Signed-off-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

```
include/linux/cgroup.h | 1 +
kernel/cgroup.c       | 7 +++++++
2 files changed, 8 insertions(+)
```

Index: devel-2.6.23-mm1/include/linux/cgroup.h

```
=====
--- devel-2.6.23-mm1.orig/include/linux/cgroup.h
+++ devel-2.6.23-mm1/include/linux/cgroup.h
@@ -233,6 +233,7 @@ int cgroup_is_descendant(const struct cg
 struct cgroup_subsys {
     struct cgroup_subsys_state *(*create)(struct cgroup_subsys *ss,
        struct cgroup *cont);
+ void (*pre_destroy)(struct cgroup_subsys *ss, struct cgroup *cont);
     void (*destroy)(struct cgroup_subsys *ss, struct cgroup *cont);
     int (*can_attach)(struct cgroup_subsys *ss,
        struct cgroup *cont, struct task_struct *tsk);
Index: devel-2.6.23-mm1/kernel/cgroup.c
```

```
=====
--- devel-2.6.23-mm1.orig/kernel/cgroup.c
+++ devel-2.6.23-mm1/kernel/cgroup.c
@@ -2158,6 +2158,13 @@ static int cgroup_rmdir(struct inode *un
     parent = cont->parent;
     root = cont->root;
     sb = root->sb;
+ /*
+  * Notify subsyses that rmdir() request comes.
+  */
+ for_each_subsys(root, ss) {
+     if ((cont->subsys[ss->subsys_id]) && ss->pre_destroy)
+         ss->pre_destroy(ss, cont);
```

+ }

```
if (cgroup_has_css_refs(cont)) {  
    mutex_unlock(&cgroup_mutex);
```

Containers mailing list

Containers@lists.linux-foundation.org

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