

---

Subject: Re: [PATCH 1/2] cgroup: Set CGRP\_RELEASABLE when adding to a cgroup

Posted by [Colin Cross](#) on Thu, 25 Nov 2010 00:11:34 GMT

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

---

On Wed, Nov 24, 2010 at 3:54 PM, Paul Menage <menage@google.com> wrote:

> On Tue, Nov 23, 2010 at 9:37 PM, Colin Cross <ccross@android.com> wrote:

>> @@ -364,12 +372,8 @@ static void \_\_put\_css\_set(struct css\_set \*cg, int taskexit)

```
>>     struct cgroup *cgrp = link->cgrp;
>>     list_del(&link->cg_link_list);
>>     list_del(&link->cgrp_link_list);
>> -     if (atomic_dec_and_test(&cgrp->count) &&
>> -         notify_on_release(cgrp)) {
>> -         if (taskexit)
>> -             set_bit(CGRP_RELEASABLE, &cgrp->flags);
>> +     if (atomic_dec_and_test(&cgrp->count))
>>         check_for_release(cgrp);
>> -     }
```

> We seem to have lost some notify\_on\_release() checks - maybe move that  
> to check\_for\_release()?

check\_for\_release immediately calls cgroup\_is\_releasable, which checks for the same bit as notify\_on\_release. There's no need for CGRP\_RELEASABLE to depend on notify\_on\_release, or to check notify\_on\_release before calling check\_for\_release.

```
>> /* Caller must verify that the css is not for root cgroup */
>> +void __css_get(struct cgroup_subsys_state *css, int count)
>> +{
>> +    atomic_add(count, &css->refcnt);
>> +    set_bit(CGRP_RELEASABLE, &css->cgroup->flags);
>> +}
>
```

> Is css\_get() the right place to be putting this? It's not clear to me  
> why a subsystem taking a refcount on a cgroup's state should render it  
> releasable when it drops that refcount.

I matched the existing behavior, \_\_css\_put sets CGRP\_RELEASABLE when refcnt goes to 0.

> Should we maybe clear the CGRP\_RELEASABLE flag right before doing the  
> userspace callback?

Actually, I think CGRP\_RELEASABLE can be dropped entirely.

check\_for\_release is only called from \_\_css\_put, cgroup\_rmdir, and \_\_put\_css\_set (or free\_css\_set\_work after my second patch). Those all imply that \_\_css\_get, get\_css\_set, or cgroup\_create have been previously called, which are the functions that set CGRP\_RELEASABLE.

---

Containers mailing list

