

---

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

Posted by [Colin Cross](#) on Fri, 03 Dec 2010 03:07:12 GMT

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

---

On Wed, Nov 24, 2010 at 4:21 PM, Paul Menage <[menage@google.com](mailto:menage@google.com)> wrote:

> On Wed, Nov 24, 2010 at 4:11 PM, Colin Cross <[ccross@android.com](mailto:ccross@android.com)> wrote:

>>>

>>> 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.

>

> OK.

>

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

>>

>

> Ah, we do appear to have had that behaviour for a while. I don't

> remember the justification for it at this point :-)

>

>> 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.

>

> Not in one case - if we create a new cgroup and try to move a thread

> into it, but the thread is exiting as we move it, we'll call

> put\_css\_set() on the new css\_set, which will drop the refcount on the

> target cgroup back to 0. We wouldn't want the auto-release

> notification to kick in in that situation, I think.

Clearing the CGRP\_RELEASABLE bit any time after the tests in  
check\_for\_release introduces a race if \_\_css\_get is called between the  
check and clearing the bit - the cgroup will have an entry, but the  
bit will not be set. Without additional locking in \_\_css\_get, I don't  
see any way to safely clear CGRP\_RELEASABLE.

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

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

---