

---

Subject: [PATCH] memcg: Do not open code accesses to res\_counter members  
Posted by [Glauber Costa](#) on Tue, 20 Mar 2012 16:53:44 GMT

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

---

We should use the accessor res\_counter\_read\_u64 for that.  
Although a purely cosmetic change is sometimes better of delayed,  
to avoid conflicting with other people's work, we are starting to  
have people touching this code as well, and reproducing the open  
code behavior because that's the standard =)

Time to fix it, then.

Signed-off-by: Glauber Costa <glommer@parallels.com>  
Cc: Johannes Weiner <hannes@cmpxchg.org>  
Cc: Michal Hocko <mhocko@suse.cz>  
Cc: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

---

```
mm/memcontrol.c | 4 +++-
1 files changed, 2 insertions(+), 2 deletions(-)
```

```
diff --git a/mm/memcontrol.c b/mm/memcontrol.c
index 87a1e21..27c1bfa 100644
--- a/mm/memcontrol.c
+++ b/mm/memcontrol.c
@@ -3708,7 +3708,7 @@ move_account:
     goto try_to_free;
     cond_resched();
     /* "ret" should also be checked to ensure all lists are empty. */
- } while (memcg->res.usage > 0 || ret);
+ } while (res_counter_read_u64(&memcg->res, RES_USAGE) > 0 || ret);
out:
    css_put(&memcg->css);
    return ret;
@@ -3723,7 +3723,7 @@ try_to_free:
    lru_add_drain_all();
    /* try to free all pages in this cgroup */
    shrink = 1;
- while (nr_retries && memcg->res.usage > 0) {
+ while (nr_retries && res_counter_read_u64(&memcg->res, RES_USAGE) > 0) {
    int progress;

    if (signal_pending(current)) {
--
1.7.7.6
```

---

Subject: Re: [PATCH] memcg: Do not open code accesses to res\_counter

members

Posted by [KAMEZAWA Hiroyuki](#) on Wed, 21 Mar 2012 01:08:29 GMT

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

---

(2012/03/21 1:53), Glauber Costa wrote:

> We should use the accessor `res_counter_read_u64` for that.  
> Although a purely cosmetic change is sometimes better of delayed,  
> to avoid conflicting with other people's work, we are starting to  
> have people touching this code as well, and reproducing the open  
> code behavior because that's the standard =)  
>  
> Time to fix it, then.  
>  
> Signed-off-by: Glauber Costa <[glommer@parallels.com](mailto:glommer@parallels.com)>  
> Cc: Johannes Weiner <[hannes@cmpxchg.org](mailto:hannes@cmpxchg.org)>  
> Cc: Michal Hocko <[mhocko@suse.cz](mailto:mhocko@suse.cz)>  
> Cc: KAMEZAWA Hiroyuki <[kamezawa.hiroyu@jp.fujitsu.com](mailto:kamezawa.hiroyu@jp.fujitsu.com)>

Acked-by: KAMEZAWA Hiroyuki <[kamezawa.hiroyu@jp.fujitsu.com](mailto:kamezawa.hiroyu@jp.fujitsu.com)>

---

Subject: Re: [PATCH] memcg: Do not open code accesses to `res_counter`  
members

Posted by [Michal Hocko](#) on Wed, 21 Mar 2012 13:26:23 GMT

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

---

On Tue 20-03-12 20:53:44, Glauber Costa wrote:

> We should use the accessor `res_counter_read_u64` for that.  
> Although a purely cosmetic change is sometimes better of delayed,  
> to avoid conflicting with other people's work, we are starting to  
> have people touching this code as well, and reproducing the open  
> code behavior because that's the standard =)  
>  
> Time to fix it, then.

Looks good to me

Acked-by: Michal Hocko <[mhocko@suse.cz](mailto:mhocko@suse.cz)>

Thanks

>  
> Signed-off-by: Glauber Costa <[glommer@parallels.com](mailto:glommer@parallels.com)>  
> Cc: Johannes Weiner <[hannes@cmpxchg.org](mailto:hannes@cmpxchg.org)>  
> Cc: Michal Hocko <[mhocko@suse.cz](mailto:mhocko@suse.cz)>  
> Cc: KAMEZAWA Hiroyuki <[kamezawa.hiroyu@jp.fujitsu.com](mailto:kamezawa.hiroyu@jp.fujitsu.com)>  
> ---  
> mm/memcontrol.c | 4 +++

```
> 1 files changed, 2 insertions(+), 2 deletions(-)
>
> diff --git a/mm/memcontrol.c b/mm/memcontrol.c
> index 87a1e21..27c1bfa 100644
> --- a/mm/memcontrol.c
> +++ b/mm/memcontrol.c
> @@ -3708,7 +3708,7 @@ move_account:
>     goto try_to_free;
>     cond_resched();
>     /* "ret" should also be checked to ensure all lists are empty. */
> - } while (memcg->res.usage > 0 || ret);
> + } while (res_counter_read_u64(&memcg->res, RES_USAGE) > 0 || ret);
> out:
>     css_put(&memcg->css);
>     return ret;
> @@ -3723,7 +3723,7 @@ try_to_free:
>     lru_add_drain_all();
>     /* try to free all pages in this cgroup */
>     shrink = 1;
> - while (nr_retries && memcg->res.usage > 0) {
> + while (nr_retries && res_counter_read_u64(&memcg->res, RES_USAGE) > 0) {
>     int progress;
>
>     if (signal_pending(current)) {
> --
> 1.7.7.6
>
> --
> To unsubscribe, send a message with 'unsubscribe linux-mm' in
> the body to majordomo@kvack.org. For more info on Linux MM,
> see: http://www.linux-mm.org/ .
> Fight unfair telecom internet charges in Canada: sign http://stopthemeteter.ca/
> Don't email: <a href="mailto:dont@kvack.org"> email@kvack.org </a>
```

--

Michal Hocko  
SUSE Labs  
SUSE LINUX s.r.o.  
Lihovarska 1060/12  
190 00 Praha 9  
Czech Republic

---

Subject: Re: [PATCH] memcg: Do not open code accesses to res\_counter members

Posted by [akpm](#) on Thu, 05 Apr 2012 23:53:35 GMT

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

On Tue, 20 Mar 2012 20:53:44 +0400

Glauber Costa <glommer@parallels.com> wrote:

```
> We should use the accessor res_counter_read_u64 for that.
> Although a purely cosmetic change is sometimes better of delayed,
> to avoid conflicting with other people's work, we are starting to
> have people touching this code as well, and reproducing the open
> code behavior because that's the standard =)
>
> ...
>
> --- a/mm/memcontrol.c
> +++ b/mm/memcontrol.c
> @@ -3708,7 +3708,7 @@ move_account:
>     goto try_to_free;
>     cond_resched();
>     /* "ret" should also be checked to ensure all lists are empty. */
> - } while (memcg->res.usage > 0 || ret);
> + } while (res_counter_read_u64(&memcg->res, RES_USAGE) > 0 || ret);
> out:
>     css_put(&memcg->css);
>     return ret;
> @@ -3723,7 +3723,7 @@ try_to_free:
>     lru_add_drain_all();
>     /* try to free all pages in this cgroup */
>     shrink = 1;
> - while (nr_retries && memcg->res.usage > 0) {
> + while (nr_retries && res_counter_read_u64(&memcg->res, RES_USAGE) > 0) {
>     int progress;
>
>
>     if (signal_pending(current)) {
```

Actually this fixes bugs on 32-bit machines. Good luck trying to demonstrate them at runtime though ;)

---