

---

Subject: Re: [RFD][PATCH] memcg: Move Usage at Task Move  
Posted by [Balbir Singh](#) on Wed, 11 Jun 2008 08:27:34 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Paul Menage wrote:

> On Thu, Jun 5, 2008 at 6:52 PM, KAMEZAWA Hiroyuki  
> <kamezawa.hiroyu@jp.fujitsu.com> wrote:  
>> Move Usage at Task Move (just an experimental for discussion)  
>> I tested this but don't think bug-free.  
>>  
>> In current memcg, when task moves to a new cg, the usage remains in the old cg.  
>> This is considered to be not good.  
>  
> Is it really such a big deal if we don't transfer the page ownerships  
> to the new cgroup? As this thread has shown, it's a fairly painful  
> operation to support. It would be good to have some concrete examples  
> of cases where this is needed.  
>  
>

I tend to agree with Paul. One of the reasons, I did not move charges is that makes migration an expensive operation. Since migration is well controlled with permissions, we assume that the node owner what he/she is doing.

>> This is a trial to move "usage" from old cg to new cg at task move.  
>> Finally, you'll see the problems we have to handle are failure and rollback.  
>>  
>> This one's Basic algorithm is  
>>  
>> 0. can\_attach() is called.  
>> 1. count movable pages by scanning page table. isolate all pages from LRU.  
>> 2. try to create enough room in new memory cgroup  
>> 3. start moving page accounting  
>> 4. putback pages to LRU.  
>> 5. can\_attach() for other cgroups are called.  
>>  
>> A case study.  
>>  
>> group\_A -> limit=1G, task\_X's usage= 800M.  
>> group\_B -> limit=1G, usage=500M.  
>>  
>> For moving task\_X from group\_A to group\_B.  
>> - group\_B should be reclaimed or have enough room.  
>>  
>> While moving task\_X from group\_A to group\_B.  
>> - group\_B's memory usage can be changed  
>> - group\_A's memory usage can be changed  
>>

>> We accounts the resouce based on pages. Then, we can't move all resource  
>> usage at once.  
>>  
>> If group\_B has no more room when we've moved 700M of task\_X to group\_B,  
>> we have to move 700M of task\_X back to group\_A. So I implemented roll-back.  
>> But other process may use up group\_A's available resource at that point.  
>>  
>> For avoiding that, preserve 800M in group\_B before moving task\_X means that  
>> task\_X can occupy 1600M of resource at moving. (So I don't do in this patch.)  
>  
> I think that pre-reserving in B would be the cleanest solution, and  
> would save the need to provide rollback.  
>  
>> 2. Don't move any usage at task move. (current implementation.)  
>> Pros.  
>> - no complication in the code.  
>> Cons.  
>> - A task's usage is chareged to wrong cgroup.  
>> - Not sure, but I believe the users don't want this.  
>  
> I'd say stick with this unless there a strong arguments in favour of  
> changing, based on concrete needs.  
>  
>> One reason is that I think a typical usage of memory controller is  
>> fork()->move->exec(). (by libcg ?) and exec() will flush the all usage.  
>  
> Exactly - this is a good reason \*not\* to implement move - because then  
> you drag all the usage of the middleware daemon into the new cgroup.  
>

Yes. The other thing is that charges will eventually fade away. Please see the  
cgroup implementation of page\_referenced() and mark\_page\_accessed(). The  
original group on memory pressure will drop pages that were left behind by a  
task that migrates. The new group will pick it up if referenced.

[snip]

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---