
Subject: Re: [PATCH 02/10] memcg: Uncharge all kmem when deleting a cgroup.
Posted by [Glauber Costa](#) on Wed, 29 Feb 2012 16:51:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 02/28/2012 09:24 PM, Suleiman Souhlal wrote:

> On Tue, Feb 28, 2012 at 11:00 AM, Glauber Costa<glommer@parallels.com> wrote:

>> On 02/27/2012 07:58 PM, Suleiman Souhlal wrote:

>>>

>>> A later patch will also use this to move the accounting to the root

>>> cgroup.

>>>

>>

>> Suleiman,

>>

>> Did you do any measurements to figure out how long does it take, average,

>> for dangling caches to go away ? Under memory pressure, let's say

>

> Unfortunately, I don't have any such measurements, other than a very artificial:

>

> # mkdir /dev/cgroup/memory/c

> # echo 1073741824> /dev/cgroup/memory/c/memory.limit_in_bytes

> # sync&& echo 3> /proc/sys/vm/drop_caches

> # echo \$\$> /dev/cgroup/memory/c/tasks

> # find /> /dev/null

> # grep '(c)' /proc/slabinfo | wc -l

> 42

> # echo \$\$> /dev/cgroup/memory/tasks

> # rmdir /dev/cgroup/memory/c

> # grep '(c)dead' /proc/slabinfo | wc -l

> 42

> # sleep 60&& sync&& for i in `seq 1 1000`; do echo 3>

> /proc/sys/vm/drop_caches ; done

> # grep '(c)dead' /proc/slabinfo | wc -l

> 6

> # sleep 60&& grep '(c)dead' /proc/slabinfo | wc -l

> 5

> # sleep 60&& grep '(c)dead' /proc/slabinfo | wc -l

> 5

>

> (Note that this is without any per-memcg shrinking patch applied. With

> shrinking, things will be a bit better, because deleting the cgroup

> will force the dentries to get shrunk.)

>

> Some of these dead caches may take a long time to go away, but we

> haven't found them to be a problem for us, so far.

>

Ok. When we start doing shrinking, however, I'd like to see a shrink

step being done before we destroy the memcg. This way we can at least reduce the number of pages lying around.
