On Wed, Feb 29, 2012 at 8:47 AM, Glauber Costa <glommer@parallels.com> wrote:
> On 02/28/2012 07:47 PM, Suleiman Souhlal wrote:
>> I hadn't considered the fact that two cgroups could have the same base
>> name.
>> I think having a name makes it a lot easier for a human to understand
>> which cgroup is using slab, so what about having both the base name of
>> the cgroup AND its css id, so that the caches are named like
>> "dentry(5:foo)"?
>
>
> That would be better, so if you really want to keep names, I'd advise you to
> go this route.
>
> However, what does "5" really means to whoever is reading that? css ids are
> not visible to the user, so there isn't really too much information you can
> extract for that. So why not only dentry-5 ?

dentry-5 gives even less information.. :-)

>> This should let us use the name of the cgroup while still being able
>> to distiguish cgroups that have the same base name.
>
>
> I am fine with name+css_id if you really want names on it.
>
>
>>> I was thinking: How about we don't bother to show them at all, and
>>> instead,
>>> show a proc-like file inside the cgroup with information about that
>>> cgroup?
>>
>>
>> One of the patches in the series adds a per-memcg memory.kmem.slabinfo.
>
> I know. What I was wondering was if we wanted to show only the non-cgroup
> slabs in /proc/slabinfo, and then show the per-cgroup slabs in the cgroup
> only.

I think /proc/slabinfo should show all the slabs in the system, to
avoid confusion.

Thanks for your comments so far.
I will try to get a v2 out soon that addresses them.

-- Suleiman