
Subject: Re: [RFC/PATCH] cgroup swap subsystem
Posted by [Pavel Emelianov](#) on Thu, 06 Mar 2008 08:38:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

KAMEZAWA Hiroyuki wrote:

> On Thu, 06 Mar 2008 11:20:17 +0300

> Pavel Emelyanov <xemul@openvz.org> wrote:

>

>> KAMEZAWA Hiroyuki wrote:

>>> On Wed, 05 Mar 2008 17:14:12 +0300

>>> Pavel Emelyanov <xemul@openvz.org> wrote:

>>>> Strongly agree. Nobody's interested in swap as such: it's just

>>>> secondary memory, where RAM is primary memory. People want to

>>>> control memory as the sum of the two; and I expect they may also

>>>> want to control primary memory (all that the current memcg does)

>>>> within that. I wonder if such nesting of limits fits easily

>>>> into cgroups or will be problematic.

>>>> This nesting would affect the res_counter abstraction, not the

>>>> cgroup infrastructure. Current design of resource counters doesn't

>>>> allow for such thing, but the extension is a couple-of-lines patch :)

>>>>

>>> IMHO, keeping res_counter simple is better.

>>>

>>> Is this kind of new entry in mem_cgroup not good ?

>>> ==

>>> struct mem_cgroup {

>>> ...

>>> struct res_counter memory_limit.

>>> struct res_counter swap_limit.

>>> ..

>>> }

>> I meant the same thing actually. By "nesting would affect" I

>> meant, that we might want to make res_counters hierarchical.

>>

>> That would kill two birds with one stone - we will make a true

>> hierarchical memory accounting and let charging of two counters

>> with one call.

>

> Hierarchical res_counter makes sense.

> Making it in simple/reasonable style will be our challenge.

I have this in my TODO list. Since this is not so urgent, then if you don't mind I can prepare the patches next week - after I set the git tree up. This change doesn't seem that big.

> Thanks,

> -Kame

>

>

Containers mailing list

Containers@lists.linux-foundation.org

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