
Subject: Re: [RFC][PATCH] another swap controller for cgroup
Posted by [Daisuke Nishimura](#) on Thu, 15 May 2008 12:01:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 2008/05/15 17:56 +0900, YAMAMOTO Takashi wrote:

>>> > If so, why is this better
>>> > than charging for actual swap usage?
>>>
>>> its behaviour is more deterministic and it uses less memory.
>>> (than nishimura-san's one, which charges for actual swap usage.)
>>>

Consuming more memory cannot be helped for my controller...

>> Using less memory is good, but maybe not worth it if the result isn't so useful.
>>
>> I'd say that it's less deterministic than nishimura-san's controller -
>> with his you just need to know how much swap is in use (which you can
>> tell by observing the app on a real system) but with yours you also
>> have to know whether there are any processes sharing anon pages (but
>> not mms).
>
> deterministic in the sense that, even when two or more processes
> from different cgroups are sharing a page, both of them, rather than
> only unlucky one, are always charged.
>

I'm not sure whether this behavior itself is good or bad,
but I think it's not good idea to make memory controller,
which charges only one process for a shared page,
and swap controller behave differently.
I think it will be confusing for users. At least,
I would feel it strange.

> another related advantage is that it's possible to move charges
> quite precisely when moving a task among cgroups.
>

Moving charges is one of future todo of my controller.
But, as you say, it won't be so precise as yours.

Thanks,
Daisuke Nishimura.

Containers mailing list

