
Subject: Re: [RFC/PATCH] cgroup swap subsystem
Posted by [Daisuke Nishimura](#) on Fri, 07 Mar 2008 08:22:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi.

kamezawa.hiroyu@jp.fujitsu.com wrote:

>>> At first look, remembering mm struct is not very good.
>>> Remembering swap controller itself is better.
>> The swap_cgroup when the page(and page_cgroup) is allocated and
>> the swap_cgroup when the page is going to be swapped out may be
>> different by swap_cgroup_move_task(), so I think swap_cgroup
>> to be charged should be determined at the point of swapout.
>>
> Accounting swap against an entity which allocs anon memory is
> not strange. Problem here is move_task itself.
> Now, charges against anon is not moved when a task which uses it
> is moved. please fix this behavior first if you think this is
> problematic.
>
> But, finally, a daemon driven by process event connector
> determines the group before process starts using anon. It's
> doubtful that it's worth to add complicated/costly ones.
>

I agree with you.

I think the current behavior of move_task is problematic,
and should fix it.
But fixing it would be difficult and add a costly process,
so I should consider more.

Thanks,
Daisuke Nishimura.

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