Subject: Re: [PATCH 6/6 mm] memcgroup: revert swap_state mods
Posted by KAMEZAWA Hiroyuki on Fri, 09 Nov 2007 09:21:56 GMT

On Fri, 9 Nov 2007 07:14:22 +0000 (GMT)
Hugh Dickins <hugh@veritas.com> wrote:

> If we're charging rss and we're charging cache, it seems obvious that
> we should be charging swapcache - as has been done. But in practice
> that doesn't work out so well: both swapin readahead and swapoff leave
> the majority of pages charged to the wrong cgroup (the cgroup that
> happened to read them in, rather than the cgroup to which they belong).
>

Thank you. I welcome this patch :)

Could I confirm a change in the logic ?

 * Before this patch, wrong swapcache charge is added to one who
   called try_to_free_page().

 * After this patch, anonymous page's charge will drop to 0 when
   page_remove_rmap() is called.

Regards,
-Kame

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