
Subject: Re: [PATCH 3/6 mm] memcgroup: fix try_to_free order

Posted by [Balbir Singh](#) on Mon, 12 Nov 2007 05:05:34 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hugh Dickins wrote:

> Why does try_to_free_mem_cgroup_pages try for order 1 pages? It's called
> when mem_cgroup_charge_common would go over the limit, and that's adding
> an order 0 page. I see no reason: it has to be a typo: fix it.

>

> Signed-off-by: Hugh Dickins <hugh@veritas.com>

> ---

> Insert just after memory-controller-add-per-container-lru-and-reclaim-v7.patch

>

> mm/vmscan.c | 2 +-

> 1 file changed, 1 insertion(+), 1 deletion(-)

>

> --- patch2/mm/vmscan.c 2007-11-08 15:46:21.000000000 +0000

> +++ patch3/mm/vmscan.c 2007-11-08 15:48:08.000000000 +0000

> @@ -1354,7 +1354,7 @@ unsigned long try_to_free_mem_cgroup_pag

> .may_swap = 1,

> .swap_cluster_max = SWAP_CLUSTER_MAX,

> .swappiness = vm_swappiness,

> - .order = 1,

> + .order = 0,

> .mem_cgroup = mem_cont,

> .isolate_pages = mem_cgroup_isolate_pages,

> };

Thanks for catching this, it is a typo

Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>

--

Warm Regards,

Balbir Singh

Linux Technology Center

IBM, ISTL

Containers mailing list

Containers@lists.linux-foundation.org

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