
Subject: Re: [RFC/PATCH] cgroup swap subsystem
Posted by [KAMEZAWA Hiroyuki](#) on Wed, 05 Mar 2008 07:01:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 05 Mar 2008 14:59:05 +0900

Daisuke Nishimura <nishimura@mfp.nes.nec.co.jp> wrote:

```
> +int swap_cgroup_charge(struct page *page,
> +  struct swap_info_struct *si,
> +  unsigned long offset)
> +{
> + int ret;
> + struct page_cgroup *pc;
> + struct mm_struct *mm;
> + struct swap_cgroup *swap;
> +
> + BUG_ON(!page);
> +
> + /*
> + * Pages to be swapped out should have been charged by memory cgroup,
> + * but very rarely, pc would be NULL (pc is not reliable without lock,
> + * so I should fix here).
> + * In such cases, we charge the init_mm now.
> + */
> + pc = page_get_page_cgroup(page);
> + if (WARN_ON(!pc))
> + mm = &init_mm;
> + else
> + mm = pc->pc_mm;
> + BUG_ON(!mm);
> +
> + rCU_read_lock();
> + swap = rCU_dereference(mm->swap_cgroup);
> + rCU_read_unlock();
> + BUG_ON(!swap);
> +
> + ret = res_counter_charge(&swap->res, PAGE_SIZE);
> + if (!ret) {
> + css_get(&swap->css);
> + si->swap_cgroup[offset] = swap;
> + }
> +
```

I think it's better to reclaim swap_entry used for SwapCache but not in Harddisk before failure.

Thanks,
-Kame

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