
Subject: Re: [PATCH 2/4] swapcgroup: add member to swap_info_struct for cgroup
Posted by [KAMEZAWA Hiroyuki](#) on Thu, 22 May 2008 09:33:35 GMT

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

On Thu, 22 May 2008 17:46:54 +0900

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

> > ==

> > #ifdef CONFIG_CGROUP_SWAP_RES_CTR

> > void swap_cgroup_init_memcg(p, memcg)

> > {

> > do something.

> > }

> > #else

> > void swap_cgroup_init_memcg(p, memcg)

> > {

> > }

> > #endif

> > ==

> >

> I think swap_cgroup_init_memcg should return old value

> of p->memcg, and I would like to name it swap_cgroup_clear_memcg,

> because it is called by sys_swapoff, so "clear" rather than "init"

> would be better.

>

> How about something like this?

>

> struct mem_cgroup **swap_cgroup_clear_memcg(p, memcg)

> {

> struct mem_cgroup **mem;

>

> mem = p->memcg;

> p->memcg = NULL;

>

> return mem;

> }

>

> and at sys_swapoff():

>

> struct mem_cgroup **memcg;

> :

> memcg = swap_cgroup_clear_memcg(p, memcg);

> :

> if (memcg)

> vfree(memcg);

>

seems good.

```
> >> + #ifdef CONFIG_CGROUP_SWAP_RES_CTLR
> >> + p->memcg = vmalloc(maxpages * sizeof(struct mem_cgroup *));
> >> + if (!p->memcg) {
> >> + error = -ENOMEM;
> >> + goto bad_swap;
> >> + }
> >> + memset(p->memcg, 0, maxpages * sizeof(struct mem_cgroup *));
> >> + #endif
> > void alloc_swap_ctlr_memcg(p)
```

>>

> OK.

> I'll implement swap_cgroup_alloc_memcg.

>

>> But this implies swapon will fail at memory shortage. Is it good ?

>>

> Hum.

> Would it be better to just disabling this feature?

>

I have no good idea. IMHO, adding printk() to show 'fatal status of not-enough-memory-for-vmalloc' will be first step.

I believe vmalloc() tend not to fail on 64bit machine, but on i386, vmalloc area is not enough.

Thanks,
-Kame

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