
Subject: Re: [PATCH -mm 5/5] swapcgroup (v3): implement force_empty
Posted by [Daisuke Nishimura](#) on Fri, 04 Jul 2008 12:33:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 4 Jul 2008 19:16:38 +0900, KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> wrote:

> On Fri, 4 Jul 2008 15:24:23 +0900

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

>

> > This patch implements force_empty of swapcgroup.

> >

> > Currently, it simply uncharges all the charges from the group.

> >

> > I think there can be other implementations.

> >

> > What I thought are:

> > - move all the charges to its parent.

> > - unuse(swap in) all the swap charged to the group.

> >

> 3. move all swap back to memory (see swapoff.)

>

>

Do you mean swapping in all the swap including used by other groups?

It would be one choice anyway.

> > But in any case, I think before implementing this way,

> > hierarchy and move charges support are needed.

> >

> > So I think this is enough for the first step.

> >

>

> I don't think hierarchy/automatic-load-balancer for swap cg is necessary.

It's the problem of how the "hierarchy" would be, I think.

I'm saying "hierarchy" here just to mean "some kind of feature where a parent includes their children".

I think "hierarchy" is needed if we implement the choice 1 above, and I personally think it would be the best choice.

> Hmm...but handling limit_change (at least, returns -EBUSY) will be necessary.

I think so too.

But I'm not sure now it's good or bad to support shrinking at limit_change about swap.

Shrinking swap means increasing the memory usage and that may cause another swapout.

> Do you consider a some magical way to move pages in swap back to memory ?

>

In this patch, I modified the `find_next_to_unuse()` to find the entry charged to a specific group.

It might be possible to modify `try_to_unuse()` (or define another function based on `try_to_unuse()`) to reduce swap usage of a specified group down to some threshold.

But, I think, one problem here is from which device the swaps should be back to memory, or usage balance between swap devices.

> In general, I like this set but we can't change the limit on demand. (maybe)

> (just putting it to TO-DO-List is okay to me.)

>

I'm sorry but what do you mean by "change the limit on demand"?

Could you explain more?

Thanks,
Daisuke Nishimura.

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