
Subject: Re: [PATCH] memory cgroup enhancements [1/5] force_empty for memory cgroup

Posted by [KAMEZAWA Hiroyuki](#) on Wed, 17 Oct 2007 05:26:42 GMT

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

On Wed, 17 Oct 2007 10:35:58 +0530

Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

> > If the only use of this is for rmdir, why not just make it part of the
> > rmdir operation on the memory cgroup if there are no tasks by default?

> >

>

> That's a good idea, but sometimes an administrator might want to force
> a cgroup empty and start fresh without necessarily deleting the cgroup.

>

I'll make a "automatic force_empty at rmdir()" patch as another patch depends on this. If we make consensus that "force_empty interface is redundant", I'll remove it later.

> I am not convinced of this hack either, specially the statement of
> setting count to SWAP_CLUSTER_MAX.

>

Just because I think there should be "unlock and rest" in this busy loop,
I need some number. Should I define other number ?

as

```
#define FORCE_RECALIM_BATCH (128)
```

```
> >> + /* drop all page_cgroup in inactive_list */
```

```
> >> + mem_cgroup_force_empty_list(mem, &mem->inactive_list);
```

```
> >> + }
```

```
> >
```

```
> > This implementation as a while loop looks very suspect since
```

```
> > mem_cgroup_force_empty_list() uses while (!list_empty(list)) as well.
```

```
> > Perhaps it's just easier here as
```

```
> >
```

```
> > if (list_empty(&mem->active_list) && list_empty(&mem->inactive_list))
```

```
> > return 0;
```

```
> >
```

```
>
```

```
> Do we VM_BUG_ON() in case the lists are not empty after calling
```

```
> mem_cgroup_force_empty_list()
```

```
>
```

Okay, I will add.

> > Reading memory.force_empty is pretty useless, so why allow it to be read
> > at all?

>
> I agree, this is not required. I wonder if we could set permissions at
> group level to mark this file as *write only*. We could use the new
> read_uint and write_uint callbacks for reading/writing integers.
>
ok, will remove.

Thanks,
-Kame

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