
Subject: Re: Supporting overcommit with the memory controller
Posted by [KAMEZAWA Hiroyuki](#) on Thu, 06 Mar 2008 09:04:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, 06 Mar 2008 11:55:47 +0300
Pavel Emelyanov <xemul@openvz.org> wrote:

> >> Can Balbir's soft-limit patches help ?
>
> [snip]
>
> >
> > Yes, that could be a useful part of the solution - I suspect we'd need
> > to have kswapd do the soft-limit push back as well as in
> > try_to_free_pages(), to avoid the high-priority jobs getting stuck in
> > the reclaim code. It would also be nice if we had:
>
> BTW, one of the way OpenVZ users determine how much memory they
> need for containers is the following: they set the limits to
> maximal values and then check the "maxheld" (i.e. the maximal level
> of consumption over the time) value.
>
> Currently, we don't have such in res_counters and I'm going to
> implement this. Objections?
>
Basically, no objection.

BTW, which does it means ?

- create a new cgroup to accounting max memory consumption, etc...

or

- add new member to mem_cgroup

or

- add new member to res_counter

Thanks,
-Kame

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