
Subject: Re: Supporting overcommit with the memory controller
Posted by [Pavel Emelianov](#) on Thu, 06 Mar 2008 09:07:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

KAMEZAWA Hiroyuki wrote:

> On Thu, 06 Mar 2008 11:55:47 +0300

> Pavel Emelianov <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

The third one - new member on res_counter. This will cost us 8 more bytes on mem_cgroup and no performance impact, since the new field is about to be touched only together with the limit and usage ones, and thus is in one cacheline.

> Thanks,

> -Kame

>

>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
