
Subject: Re: [RFC][PATCH] memory controller per zone patches take 2 [9/10]
per-zone-lru for memory cgroup

Posted by [Balbir Singh](#) on Mon, 19 Nov 2007 08:34:23 GMT

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

KAMEZAWA Hiroyuki wrote:

> On Mon, 19 Nov 2007 11:51:15 +0530

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

>>> =

>>> /cgroup/group_A/group_A_1

>>> . /group_A_2

>>> /group_A_3

>>> (LRU(s) will be used for maintaining parent/child groups.)

>>>

>> The LRU's will be shared, my vision is

>>

>> LRU

>> ^ ^

>> ||

>> Mem-----+ +----Mem

>>

>>

>> That two or more mem_cgroup's can refer to the same LRU list and have

>> their own resource counters. This setup will be used in the case

>> of a hierarchy, so that a child can share memory with its parent

>> and have it's own limit.

>>

>> The mem_cgroup will basically then only contain a reference

>> to the LRU list.

>>

> Hmm, interesting.

>

> Then,

> group_A_1's usage + group_A_2's usage + group_A_3's usgae < group_A's limit.

> group_A_1, group_A_2, group_A_3 has its own limit.

Yes that is correct

> In plan.

>

> I wonder if we want rich control functions, we need "share" or "priority" among

> childs. (but maybe this will be complicated one.)

>

That would nice and the end goal of providing this feature. We also need
to provide soft-limits (more complex) and guarantees (with the kernel
memory controller coming in - nice to have, but not necessary for now)

> Thank you for explanation.
>
> Regards,
> -Kame
>

Thanks for helping out the memory controller.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

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