
Subject: Re: [PATCH v6 3/8] foundations of per-cgroup memory pressure controlling.

Posted by [Glauber Costa](#) on Thu, 13 Oct 2011 08:25:53 GMT

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

On 10/13/2011 09:53 AM, KAMEZAWA Hiroyuki wrote:

> On Mon, 10 Oct 2011 14:24:23 +0400

> Glauber Costa<glommer@parallels.com> wrote:

>

>> This patch converts struct sock fields memory_pressure,
>> memory_allocated, sockets_allocated, and sysctl_mem (now prot_mem)
>> to function pointers, receiving a struct mem_cgroup parameter.

>>

>> enter_memory_pressure is kept the same, since all its callers
>> have socket a context, and the kmem_cgroup can be derived from
>> the socket itself.

>>

>> To keep things working, the patch convert all users of those fields
>> to use accessor functions.

>>

>> In my benchmarks I didn't see a significant performance difference
>> with this patch applied compared to a baseline (around 1 % diff, thus
>> inside error margin).

>>

>> Signed-off-by: Glauber Costa<glommer@parallels.com>

>> CC: David S. Miller<davem@davemloft.net>

>> CC: Hiroyouki Kamezawa<kamezawa.hiroyu@jp.fujitsu.com>

>> CC: Eric W. Biederman<ebiederm@xmission.com>

>

> Reviewed-by: KAMEZAWA Hiroyuki<kamezawa.hiroyu@jp.fujitsu.com>

>

> a nitpick.

>

>

>> #ifdef CONFIG_INET

>> +enum {

>> + UNDER_LIMIT,

>> + OVER_LIMIT,

>> +};

>> +

>

> It may be better to move this to res_counter.h or memcontrol.h

>

Sorry Kame,

It is in memcontrol.h already. What exactly do you mean here ?
