
Subject: Re: Re: [PATCH v5 00/10] per-cgroup tcp memory pressure
Posted by [KAMEZAWA Hiroyuki](#) on Tue, 22 Nov 2011 02:07:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 18 Nov 2011 17:39:03 -0200

Glauber Costa <glommer@parallels.com> wrote:

> On 11/17/2011 07:35 PM, David Miller wrote:

> > TCP specific stuff in mm/memcontrol.c, at best that's not nice at all.

>

> How crucial is that? Thing is that as far as I am concerned, all the

> memcg people really want the inner layout of struct mem_cgroup to be

> private to memcontrol.c

This is just because memcg is just related to memory management and I don't want it be wide spreaded, 'struct mem_cgroup' has been changed often.

But I don't like to have TCP code in memcgroup.c.

New idea is welcome.

> This means that at some point, we need to have

> at least a wrapper in memcontrol.c that is able to calculate the offset

> of the tcp structure, and since most functions are actually quite

> simple, that would just make us do more function calls.

>

> Well, an alternative to that would be to use a void pointer in the newly

> added struct cg_proto to an already parsed memcg-related field

> (in this case tcp_memcontrol), that would be passed to the functions

> instead of the whole memcg structure. Do you think this would be

> preferable ?

>

like this ?

```
struct mem_cgroup_sub_controls {
    struct mem_cgroup *mem;
    union {
        struct tcp_mem_control tcp;
    } data;
};
/* for loosely coupled controls for memcg */
struct memcg_sub_controls_function
{
    struct memcg_sub_controls (*create)(struct mem_cgroup *);
    struct memcg_sub_controls (*destroy)(struct mem_cgroup *);
}
```

```
int register_memcg_sub_controls(char *name,
```

```
struct memcg_sub_controls_function *abis);
```

```
struct mem_cgroup {  
    .....  
    .....  
    /* Root memcg will have no sub_controls! */  
    struct memcg_sub_controls *sub_controls[NR_MEMCG_SUB_CONTROLS];  
}
```

Maybe some functions should be exported.

Thanks,
-Kame
