
Subject: [PATCH v2] cgroup: prefer [kv]zalloc[_node] over [kv]malloc+memset in memory controller code.

Posted by [Jesper Juhl](#) on Thu, 04 Nov 2010 20:17:27 GMT

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

Hi,

In `mem_cgroup_alloc()` we currently do either `kmalloc()` or `vmalloc()` then followed by `memset()` to zero the memory. This can be more efficiently achieved by using `kzalloc()` and `vzalloc()`.

There's also one situation where we can use `kzalloc_node()` - this is what's new in this version of the patch.

The original patch was:

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

Reviewed-by: Minchan Kim <minchan.kim@gmail.com>

Reviewed-by: Wu Fengguang <fengguang.wu@intel.com>

Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>

Here's version 2. I'd appreciate it if someone could merge it, but I don't know who that someone would be.

Signed-off-by: Jesper Juhl <jj@chaosbits.net>

memcontrol.c | 9 +++-----

1 file changed, 3 insertions(+), 6 deletions(-)

diff --git a/mm/memcontrol.c b/mm/memcontrol.c

index 9a99cfa..4f4e676 100644

--- a/mm/memcontrol.c

+++ b/mm/memcontrol.c

@@ -4169,13 +4169,11 @@ static int alloc_mem_cgroup_per_zone_info(struct mem_cgroup
*mem, int node)

*/

if (!node_state(node, N_NORMAL_MEMORY))

tmp = -1;

- pn = kmalloc_node(sizeof(*pn), GFP_KERNEL, tmp);

+ pn = kzalloc_node(sizeof(*pn), GFP_KERNEL, tmp);

if (!pn)

return 1;

mem->info.nodeinfo[node] = pn;

- memset(pn, 0, sizeof(*pn));

-

for (zone = 0; zone < MAX_NR_ZONES; zone++) {

mz = &pn->zoneinfo[zone];

```

for_each_lru(l)
@@ -4199,14 +4197,13 @@ static struct mem_cgroup *mem_cgroup_alloc(void)

/* Can be very big if MAX_NUMNODES is very big */
if (size < PAGE_SIZE)
- mem = kmalloc(size, GFP_KERNEL);
+ mem = kzalloc(size, GFP_KERNEL);
else
- mem = vmalloc(size);
+ mem = vzalloc(size);

if (!mem)
return NULL;

- memset(mem, 0, size);
mem->stat = alloc_percpu(struct mem_cgroup_stat_cpu);
if (!mem->stat) {
if (size < PAGE_SIZE)

```

--

Jesper Juhl <jj@chaosbits.net> <http://www.chaosbits.net/>
Plain text mails only, please <http://www.expita.com/nomime.html>
Don't top-post <http://www.catb.org/~esr/jargon/html/T/top-post.html>

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

Subject: Re: [PATCH v2] cgroup: prefer [kv]zalloc[_node] over [kv]malloc+memset
in memory controller code.

Posted by [Li Zefan](#) on Thu, 04 Nov 2010 20:35:03 GMT

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

```

> In mem_cgroup_alloc() we currently do either kmalloc() or vmalloc() then
> followed by memset() to zero the memory. This can be more efficiently
> achieved by using kzalloc() and vzalloc().
> There's also one situation where we can use kzalloc_node() - this is
> what's new in this version of the patch.
>
> The original patch was:
>
> Acked-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
> Reviewed-by: Minchan Kim <minchan.kim@gmail.com>
> Reviewed-by: Wu Fengguang <fengguang.wu@intel.com>
> Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>

```

>
> Here's version 2. I'd appreciate it if someone could merge it, but I don't
> know who that someone would be.
>

Normally it's Andrew Morton.

btw, a better title is: [...] memcggroup: prefer ... over ... memset

Containers mailing list

Containers@lists.linux-foundation.org

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