

---

Subject: Re: [PATCH -mm 1/3] cgroup: use a hash table for css\_set finding  
Posted by [Li Zefan](#) on Thu, 03 Apr 2008 07:51:05 GMT

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

---

KAMEZAWA Hiroyuki wrote:

```
> On Thu, 03 Apr 2008 13:52:43 +0800
> Li Zefan <lizf@cn.fujitsu.com> wrote:
>> +/* hash table for cgroup groups. This improves the performance to
>> + * find an existing css_set */
>> +#define CSS_SET_HASH_BITS 7
>> +#define CSS_SET_TABLE_SIZE (1 << CSS_SET_HASH_BITS)
>> +static struct hlist_head css_set_table[CSS_SET_TABLE_SIZE];
>
> How above number is selected ?
>
```

I suppose 100 will be suitable, so i would like to choose from 6 or 7 bits.

```
>> +static struct hlist_head *css_set_hash(struct cgroup_subsys_state *css[])
>> +{
>> + int i;
>> + int index;
>> + unsigned long tmp = 0UL;
>> +
>> + for (i = 0; i < CGROUP_SUBSYS_COUNT; i++)
>> + tmp += (unsigned long)css[i];
>> +
>
> maybe css[i]'s lower 2-3 bits will be ignored. because they are always 0.
>
> And I don't like "+" for hash. how about
> ==
> for (i = 0; i < CGROUP_SUBSYS_COUNT; i++)
> unsigned long x;
> x = (unsigned long)css[i];
> tmp = (x >> 16) ^ (x & 0xffff)
> ==
> or some func, which uses full bits.
>
```

I'm using hash\_long(), which has been proved to be a good one. And I've tested css\_set\_hash(), I run the css\_set benchmark with N == 1000, the for loop in find\_existing\_css\_set() will break out within 10 iterations for most cases, which is the expected result.

```
>
>> + index = hash_long(tmp, CSS_SET_HASH_BITS);
>> +
```

```
>> + return &css_set_table[index];
>> +}
>> +
>> /* We don't maintain the lists running through each css_set to its
>> * task until after the first call to cgroup_iter_start(). This
>> * reduces the fork()/exit() overhead for people who have cgroups
>> @@ -219,6 +240,7 @@ static int use_task_css_set_links;
>> static void unlink_css_set(struct css_set *cg)
>> {
>> write_lock(&css_set_lock);
>> + hlist_del(&cg->hlist);
>> list_del(&cg->list);
>> css_set_count--;
>
> This css_set_lock is worth to be rwlock ?
> how about per hashline spinlock ? (but per-hashline seems overkill..)
>
```

I think it's an overkill. And the `css_set_lock` protects not only the hash table.

Thanks for looking into this. :)

Regards,  
Li Zefan

---

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

---