
Subject: Re: [RFC][mm] [1/2] Simple stats for cpu resource controller

Posted by [Balbir Singh](#) on Sun, 06 Apr 2008 05:12:17 GMT

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

Balaji Rao wrote:

```
> On Sunday 06 April 2008 02:29:14 am Dhaval Giani wrote:
>> On Sun, Apr 06, 2008 at 02:01:52AM +0530, Balaji Rao wrote:
>>> On Sunday 06 April 2008 01:10:41 am Dhaval Giani wrote:
>>>> +};
>>>> +
>>>> +struct cpu_cgroup_stat_cpu {
>>>> + s64 count[CPU_CGROUP_STAT_NSTATS];
>>>> u64? time does not go negative :)
>>> Right. But these stats are not only going to measure time. We need the
> same
>>> variables for measuring other stats as well. I'm not sure if we would
>>> encounter scheduler stats that would count negative.
>>>
>>> Balbir, what do you say ?
>> I would prefer to keep the stats logically separate. So something like
>> struct cpu_cgroup_stat_cpu {
>>   u64 time[];
>>   s64 some_other_stat;
>> }
>> and so on. (I am not sure, is there some advantage gained by using
>> structs?) Makes the code more maintainable imho.
>>
> This would break the generic nature of __cpu_cgroup_stat_add. Its not a nice
> thing in my opinion.
>
```

I prefer keeping stats in the array as Balaji has done, it makes it easier to do batch processing on the stats.

--

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>
