

---

Subject: Re: [RFC][mm] Simple stats for cpu resource controller v3  
Posted by [akpm](#) on Sat, 03 May 2008 00:19:38 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Sat, 3 May 2008 05:26:46 +0530  
Balaji Rao <balajirao@gmail.com> wrote:

> > yes, that would be good.  
> OK, so when does account\_system\_time get called for the first time ? after  
> IRQs are set up, is it ? So, where do we place the hook ?

Don't know - I'd need to dive in and work that out, and it's probably  
better than you do this..

> Here's the patch.  
>  
> diff --git a/include/linux/percpu\_counter.h b/include/linux/percpu\_counter.h  
> index 9007ccd..8a1b756 100644  
> --- a/include/linux/percpu\_counter.h  
> +++ b/include/linux/percpu\_counter.h  
> @@ -21,7 +21,7 @@ struct percpu\_counter {  
> #ifdef CONFIG\_HOTPLUG\_CPU  
> struct list\_head list; /\* All percpu\_counters are on a list \*/  
> #endif  
> - s32 \*counters;  
> + s32 counters[NR\_CPUS];  
> };

Please, no. That's a 4092-byte increase in sizeof(struct percpu\_counter).  
Hence a 12 kbyte increase in sizeof(struct ext3\_sb\_info). Let's just sort  
out the cgroup startup ordering.

<looks at \_\_percpu\_alloc\_mask>  
<wanders off-topic>

Eric, is that optimal? alloc\_percpu() will pass down cpu\_possible\_map in  
'mask', and we only need to allocate enough slots to cover the  
highest-set-bit in cpu\_possible\_map. However the implementation ignores  
'mask' and does

```
size_t sz = roundup(nr_cpu_ids * sizeof(void *), cache_line_size());  
void *pdata = kzalloc(sz, gfp);
```

Now, if the highest-set-bit in cpu\_possible\_map is always equal to  
(1<nr\_cpu\_ids) then it doesn't matter. But is that the case?

(If someone calls `__percpu_alloc_mask` with something that has less bits set than `cpu_possible_map` then it surely is wasteful, but that sounds unlikely).

---

Containers mailing list

Containers@lists.linux-foundation.org

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

---