
Subject: Re: Per container statistics (containerstats)
Posted by [Andrew Morton](#) on Fri, 08 Jun 2007 02:39:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 08 Jun 2007 07:51:12 +0530 Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

> Andrew Morton wrote:
> >
> > I'd have hoped to see containerstats.c in here.
> >
>
> The current statistics code is really small, so it fit into taskstats.c.
> May be in the future, we could re-factor it and move it out.

I was referring to your userspace tool which reads this stuff. The one which you described in the changelog.

```
> >> + rcu_read_lock();  
> >> +  
> >> + for_each_root(root) {  
> >> + if (!root->subsys_bits)  
> >> + continue;  
> >> + root_cont = &root->top_container;  
> >> + get_first_subsys(root_cont, NULL, &subsys_id);  
> >> + do_each_thread(g, p) {  
> >  
> > this needs tasklist_lock?  
> >  
>  
> rcu_read_lock() should be fine. From Eric's patch at  
>  
> 2.6.17-mm2 - proc-remove-tasklist_lock-from-proc_pid_readdir.patch  
>  
> The patch mentions that "We don't need the tasklist_lock to safely  
> iterate through processes anymore."  
>
```

oh, OK. rcu_read_lock() is the new lock_kernel() - always hard to tell what it's locking.
