
Subject: Re: [PATCH 2/7] containers (V7): Cpusets hooked into containers
Posted by [Paul Menage](#) on Wed, 07 Mar 2007 16:12:20 GMT

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

On 3/7/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:

> On Mon, Feb 12, 2007 at 12:15:23AM -0800, menage@google.com wrote:

```
> > - mutex_lock(&callback_mutex);  
> > - list_add(&cs->sibling, &cs->parent->children);  
> > + cont->cpuset = cs;  
> > + cs->container = cont;  
> >   number_of_cpusets++;  
> > - mutex_unlock(&callback_mutex);  
>
```

> What's the rule to read/write number_of_cpusets? The earlier cpuset code was
> incrementing/decrementing under callback_mutex, but now we aren't. How safe is
> that?

We're still inside manage_mutex, so we guarantee that no-one else is
changing it.

>
> The earlier cpuset code also was reading number_of_cpusets w/o the
> callback_mutex held (atleast in cpuset_zone_allowed_softwall). Is that safe?

Yes, I think so. Unless every memory allocator was to hold a lock for
the duration of alloc_pages(), number_of_cpusets can theoretically be
out of date by the time you're using it. But since the process could
have allocated just before you created the first cpuset and moved it
into that cpuset anywa, it's not really a race (and the consequences
are inconsequential).

Paul
