

Paul Menage wrote:

> On 5/1/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:
>> For the CPU controller I was working on, (a fast access to) such a list would
>> have been valuable. Basically each task has a weight associated with it
>> (p->load_weight) which is made to depend upon its class limit. Whenever
>> the class limit changes, we need to go and change all its member task's
>> ->load_weight value.
>>
>> If you don't maintain the per-container task list, I guess I could still
>> work around it, by either:
>>
>> - Walk the task table and find relevant members
>
> That doesn't seem like a terrible solution to me, unless you expect
> the class limit to be changing incredibly frequently.
>
> If we had multiple subsystems that needed to walk the container member
> list on a fast-path operation (e.g. to make a scheduling decision)
> that would be a good reason to maintain such a list.
>

I now have a use case for maintaining a per-container task list.
I am trying to build a per-container stats similar to taskstats.
I intend to support container accounting of

1. Tasks running
2. Tasks stopped
3. Tasks un-interruptible
4. Tasks blocked on IO
5. Tasks sleeping

This would provide statistics similar to the patch that Pavel had sent out.

I faced the following problems while trying to implement this feature

1. There is no easy way to get a list of all tasks belonging to a container
(we need to walk all threads)
2. There is no concept of a container identifier. When a user issues a command
to extract statistics, the only unique container identifier is the container
path, which means that we need to do a path lookup to determine the dentry
for the container (which gets quite ugly with all the string manipulation)

Adding a container id, will make it easier to find a container and return
statistics belonging to the container.

If we get these two features added to the current patches, it will make the infrastructure more "developer" and eventually more user friendly :-)

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL
