Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Paul Jackson on Sat, 03 Mar 2007 10:21:00 GMT
View Forum Message <> Reply to Message

> Regarding semantics, can you be more specific?

Unfortunately not - sorry.

I've been off in other areas, and not found the time
to read through this current PATCH or think about it
carefully enough to be really useful.

Your reply seemed reasonable enough.

> It should have the same perf overhead as the original container patches
> (basically a double dereference - task->containers/nsproxy->cpuset -
> required to get to the cpuset from a task).

There is just one spot that this might matter to cpusets.

Except for one hook, cpusets uses the mems_allowed and cpus_allowed
masks in the task struct to avoid having to look at the cpuset on hot
code paths.

There is one RCU guarded reference per memory allocation to
 current->cpuset->mems_generation
in the call to cpuset_update_task_memory_state(), for tasks that are in
some cpuset -other- than the default top cpuset, on systems that have
explicitly created additional (other than the top cpuset) cpusets after
boot.

If that RCU guarded reference turned into taking a global lock, or
pulling in a cache line that was frequently off dirty in some other
node, that would be unfortunate.

But that's the key hook so far as cpuset performance impact is
concerned.

Perhaps you could summarize what becomes of this hook, in this
brave new world of rcfs ...

--
                  I won't rest till it's the best ...
                  Programmer, Linux Scalability
                  Paul Jackson <pj@sgi.com> 1.925.600.0401

____
Containers mailing list

Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers