Subject: Re: [PATCH] task containersv11 add tasks file interface fix for cpusets Posted by Paul Jackson on Wed, 03 Oct 2007 17:58:17 GMT

View Forum Message <> Reply to Message

Paul M wrote:

- > > The code in kernel/cgroup.c attach_task() which skips the
- > > attachment of a task to the group it is already in has to be
- > > removed. Cpusets depends on reattaching a task to its current
- >> cpuset, in order to trigger updating the cpus allowed mask in the
- > > task struct.

>

- > Can you explain a bit more about why this is needed? (i.e. specific
- > scenarios where cpusets will break without this change).

>

> What triggers the reattach in normal use? Something from userspace?

Yes, something in user space has to do it. It's part of the kernel-user cpuset API. If you change a cpuset's 'cpus', then you have to rewrite each pid in its 'tasks' file back to that 'tasks' file in order to get that 'cpus' change to be applied to the task struct cpus_allowed of each task, and thereby visible to the scheduler.

--

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.linux-foundation.org

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