

Glauber Costa wrote:

> Hello Tejun,
>
> During your cgroup refactor, I was wondering if you have any plans to get rid of the bind() callback that is called when hierarchies are moved?

I planned to remove it long ago, and Paul M agreed. But after some time, I was trying to make use of it in a patchset, which was used to fix the problem that remount with different subsys bits will fail for !root cgroups.

>
> At least in tree, there seems to be no users for that.
> I actually planned to use it myself, to start or remove a jump label
> when cpuacct and cpu cgroups were comounted.
>
> Problem is, because we have some calls in the cpuset cgroup from inside the cpu hotplug handler, we end up taking the almighty cgroup_mutex from inside the cpu_hotplug.lock.
>
> jump labels take it in most arches through the get_online_cpus() function call. This means we effectively can't apply jump labels with the cgroup_mutex held, which is the case throughout the whole bind() call.
>
> All that explained, I figured I might as well ask before I attempted a solution to that myself: as much as populate(), bind seems to be one of the overly complicated callbacks, designed for a scenario in which everything can come and go at will, which is something we're trying to fix.
>

As we aim for single hierarchy, it definately should be removed.
