Subject: Re: [PATCH 4/5] Setup the control group Posted by Paul Menage on Mon, 01 Oct 2007 16:04:12 GMT View Forum Message <> Reply to Message

On 10/1/07, Balbir Singh > balbir@linux.vnet.ibm.com wrote:

>

- > Excellent, I prefer the later as well, but it would mean overheads
- > for controllers not using the hierarchy.

I don't think it would have to with the ideas I've been thinking about - each task would still have a set of pointers to subsystems which could be dereferenced just as quickly. The complexity comes in trying to map a task to its actual cgroup object in a given hierarchy - this would involve a bit more work on the part of the cgroup framework, but wouldn't be a fast path operation.

See my mail last week titled "Thoughts on virtualizing task containers".

- > a design such that parents<->children can effectively share resources,
- > track them and do so recursively, that would be really nice.

I think the recursive tracking would probably need to be supplied by the subsystem rather than by the framework. But there's no reason that multiple subsystems couldn't re-use the same hierarchy code via e.g. resource counters. So when you initialize a resource counter you'd tell it about its parent resource counter, and it would handle the recursion automatically in charge/uncharge.

Paul