

On 09/26/2012 11:34 PM, Tejun Heo wrote:

> Hello,

>

> On Wed, Sep 26, 2012 at 10:56:09PM +0400, Glauber Costa wrote:

>> For me, it is the other way around: it makes perfect sense to have a

>> per-subtree selection of features where it doesn't hurt us, provided it

>> is hierarchical. For the mere fact that not every application is

>> interested in this (Michal is the one that is being so far more vocal

>> about this not being needed in some use cases), and it is perfectly

>> valid to imagine such applications would coexist.

>>

>> So given the flexibility it brings, the real question is, as I said,

>> backwards: what is it necessary to make it a global switch ?

>

> Because it hurts my head and it's better to keep things simple. We're

> planning to retire `use_hierarchy` in sub hierarchies and I'd really

> like to prevent another fiasco like that unless there absolutely is no

> way around it. Flexibility where necessary is fine but let's please

> try our best to avoid over-designing things. We've been far too good

> at getting lost in flexibility maze. Michal, care to chime in?

>

I would very much like to hear Michal here as well, sure.

But as I said in this very beginning of this, you pretty much know that I am heavily involved in trying to get rid of `use_hierarchy`, and by no means I consider this en pair with that.

`use_hierarchy` is a hack around a core property of cgroups, the fact that they are hierarchical. Its mere existence came to be to overcome a performance limitation.

It puts you in contradictory situation where you have cgroups organized as directories, and then not satisfied in making this hierarchical representation be gravely ignored, forces you to use nonsensical terms like "flat hierarchy", making us grasp at how it is to be a politician once in our lifetimes.

Besides not being part of cgroup core, and respecting very much both cgroups' and basic sanity properties, kmem is an actual feature that some people want, and some people don't. There is no reason to believe that applications that want will live in the same environment with ones that don't want.