
Subject: Re: [PATCH 1/2] rcfs core patch
Posted by [Paul Jackson](#) on Fri, 09 Mar 2007 09:38:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill, responding to Herbert:

> > do we need or even want that? IMHO the hierarchical
> > concept CKRM was designed with, was also the reason
> > for it being slow, unuseable and complicated
> 1. cpusets are hierarchical already. So hierarchy is required.

I think that CKRM has a harder time doing a hierarchy than cpusets.

CKRM is trying to account for and control how much of an amorphous resource is used, whereas cpusets is trying to control whether a specifically identifiable resource is used, or not used, not how much of it is used.

A child cpuset gets configured to allow certain CPUs and Nodes, and then does not need to dynamically pass back any information about what is actually used - it's a one-way control with no feedback. That's a relatively easier problem.

CKRM (as I recall it, from long ago ...) has to track the amount of usage dynamically, across parent and child groups (whatever they were called.) That's a harder problem.

So, yes, as Kirill observes, we need the hierarchy because cpusets has it, cpuset users make good use of the hierarchy, and the hierarchy works fine in that case, even if a hierarchy is more difficult for CKRM.

--

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>
