Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction
Posted by Paul Jackson on Wed, 20 Sep 2006 20:27:34 GMT
View Forum Message <> Reply to Message

Chandra wrote:
> We had this discussion more than 18 months back and concluded that it is
> not the right thing to do. Here is the link to the thread:

Because it is easy enough to carve memory up into nice little nameable
chunks, it might be the case that we can manage the percentage of
memory used by the expedient of something like cpusets and fake nodes.

Indeed, that seems to be doable, based on this latest work of Andrew
and others (David, some_bright_spark@jp, Magnus, ...). There are
still a bunch of wrinkles that remain to be ironed out.

For other resources, such as CPU cycles and network bandwidth, unless
another bright spark comes up with an insight, I don't see how to
express the "percentage used" semantics provided by something such
as CKRM, using anything resembling cpusets.

... Can one imagine having the scheduler subdivide each second of
time available on a CPU into several fake-CPUs, each one of which
speaks for one of those sub-second fake-CPU slices? Sounds too
weird to me, and a bit too rigid to be a servicable CKRM substitute.

--

I won't rest till it's the best ...
Programmer, Linux Scalability
Paul Jackson <pj@sgi.com> 1.925.600.0401