

---

Subject: Re: [ckrm-tech] [PATCH 00/10] Containers(V10): Generic Process Containers

Posted by [Paul Jackson](#) on Fri, 08 Jun 2007 20:05:17 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Serge wrote:

> Paul (Jackson), is this comment added in cpusets close enough to what  
> you were asking for?

This comment?

- + \* Currently we refuse to set up the container - thereby
- + \* refusing the task to be entered, and as a result refusing
- + \* the sys\_unshare() or clone() which initiated it - if any
- + \* sibling cpusets have exclusive cpus or mem.
- + \*
- + \* If this becomes a problem for some users who wish to
- + \* allow that scenario, then cpuset\_post\_clone() could be
- + \* changed to grant parent->cpus\_allowed-sibling\_cpus\_exclusive
- + \* (and likewise for mems) to the new container.

Nice - thanks.

--

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

---