
Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction

Posted by [Paul Jackson](#) on Thu, 21 Sep 2006 22:48:04 GMT

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

Paul M wrote:

> Page allocation and task scheduling are resource controller issues,
> not generic process container issues.

But when a process is moved to a different container, its page allocation and task scheduling constraints and metrics move too.

One of the essential differences, for example, between the two memory constraint mechanisms we have now, mempolicy.c and cpuset.c, is that mempolicy only affects the current task, so has an easier time of the locking and its hooks in the page allocation code path, whereas cpusets allows any task to change any other tasks memory constraints.

This made the cpuset hooks in the page allocation code path more difficult -- and as you have recently shown, we aren't done working that code path yet ;).

This is likely true in general for resource controllers. One of their more challenging design aspects are the hooks they require in the code paths that handle the various controlled resources.

One has to use these hooks to access the container on these fairly hot code paths. And since the container can be changing in parallel at the same time, it can be challenging to handling the necessary locking without forcing a system-wide lock there.

Doable, I presume. But challenging.

--

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