Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction
Posted by Paul Jackson on Thu, 21 Sep 2006 21:59:46 GMT
View Forum Message <> Reply to Message

Paul wrote:
> But, there's no reason that the OpenVZ resource control mechanisms
> couldn't be hooked into a generic process container mechanism along
> with cpusets and RG.

Can the generic container avoid performance bottlenecks due to locks
or other hot cache lines on the main code paths for fork, exit, page
allocation and task scheduling?

--

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

Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction
Posted by Paul Menage on Thu, 21 Sep 2006 22:07:42 GMT
View Forum Message <> Reply to Message

On 9/21/06, Paul Jackson <pj@sgi.com> wrote:
>
> Can the generic container avoid performance bottlenecks due to locks
> or other hot cache lines on the main code paths for fork, exit, page
> allocation and task scheduling?

Page allocation and task scheduling are resource controller issues,
not generic process container issues. The generic process containers
would have essentially the same overheads for fork/exit that cpusets
have currently.

Paul