Subject: Re: [ckrm-tech] [Lhms-devel] [patch00/05]: Containers(V2)- Introduction
Posted by Paul Menage on Thu, 21 Sep 2006 01:36:01 GMT

View Forum Message <> Reply to Message

On 9/20/06, KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> wrote:
> For example
>
> In following scenario,
> ==
> (1). add <pid> > /mnt/configfs/containers/my_container/add_task
> (2). <pid> does some work.
> (3). echo <pid> > /mnt/configfs/containers/my_container/rm_task
> (4). echo <pid> > /mnt/configfs/containers/my_container2/add_task
> ==
> (if fake-pgdat/memory-hotplug is used)
> The pages used by <pid> in (2) will be accounted in 'my_container' after (3).
> Is this user's wrong use of system ?

Yes. You can't use memory node partitioning for file pages in this way
unless you have strict controls over who can access the data sets in
question, and are careful to prevent people from moving between
containers. So it's not suitable for all uses of resource-isolating
containers.

Who is to say that the pages allocated in (2) *should* be reaccounted
to my_container2 after (3)? Some people might want that, other
(including me) might not.

Paul