
Subject: Re: [RFC][PATCH 4/7] UBC: syscalls (user interface)
Posted by [Andrew Morton](#) on Thu, 17 Aug 2006 15:40:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, 17 Aug 2006 16:13:30 +0400
Kirill Korotaev <dev@sw.ru> wrote:

> > I was more thinking about (for example) user land physical memory limit
> > for that bean counter. If the limits are going down, then the system
> > call should try to flush out page cache pages or swap out anonymous
> > memory. But you are right that it won't be possible in all cases, like
> > for in kernel memory limits.
> Such kind of memory management is less efficient than the one
> making decisions based on global shortages and global LRU algorithm.

I also was quite surprised that openvz appears to have no way of
constraining a container's memory usage. "I want to run this bunch of
processes in a 4.5GB container".

> The problem here is that doing swap out takes more expensive disk I/O
> influencing other users.

A well-set-up container would presumably be working against its own
spindle(s). If the operator has gone to all the trouble of isolating a job
from the system's other jobs, he'd be pretty dumb to go and let all the
"isolated" jobs share a stinky-slow resource like a disk.

But yes, swap is a problem. To do this properly we'd need a way of saying
"this container here uses that swap device over there".
