
Subject: Re: [RFC][PATCH 4/7] UBC: syscalls (user interface)

Posted by [Rohit Seth](#) on Thu, 17 Aug 2006 17:08:19 GMT

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

On Thu, 2006-08-17 at 16:13 +0400, Kirill Korotaev wrote:

> Rohit Seth wrote:

> > On Wed, 2006-08-16 at 20:04 +0100, Alan Cox wrote:

> >

> >>Ar Mer, 2006-08-16 am 11:17 -0700, ysgrifennodd Rohit Seth:

> >>

> >>>I think there should be a check here for seeing if the new limits are

> >>>lower than the current usage of a resource. If so then take appropriate

> >>>action.

> >>

> >>Generally speaking there isn't a sane appropriate action because the

> >>resources can't just be yanked.

> >>

> >

> >

> > 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.

>

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

> influencing other users.

>

> So throttling algorithms if wanted should be optional, not mandatory.

> Lets postpone it and concentrate on the core.

>

I'm really interested in seeing what changes you make in alloc_page when the container limits are hit.

When a container is throttling then yes it will have some additional cost to other containers but that is the cost of sharing an underlying platform.

-rohit
