
Subject: Re: [ckrm-tech] [PATCH 4/7] UBC: syscalls (user interface)
Posted by [Chandra Seetharaman](#) on Tue, 22 Aug 2006 18:34:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 2006-08-22 at 12:58 +0900, Magnus Damm wrote:

> On Mon, 2006-08-21 at 18:16 -0700, Rohit Seth wrote:

> > On Mon, 2006-08-21 at 11:47 +0900, Magnus Damm wrote:

> > > On Fri, 2006-08-18 at 07:45 -0700, Dave Hansen wrote:

> > > > On Fri, 2006-08-18 at 12:08 +0400, Andrey Savochkin wrote:

> > > > >

> > > > > A) Have separate memory management for each container,

> > > > > with separate buddy allocator, lru lists, page replacement mechanism.

> > > > > That implies a considerable overhead, and the main challenge there

> > > > > is sharing of pages between these separate memory managers.

> > > >

> > > > Hold on here for just a sec...

> > > >

> > > > It is quite possible to do memory management aimed at one container

> > > > while that container's memory still participates in the main VM.

> > > >

> > > > There is overhead here, as the LRU scanning mechanisms get less

> > > > efficient, but I'd rather pay a penalty at LRU scanning time than divide

> > > > up the VM, or coarsely start failing allocations.

> > >

> > > This could of course be solved with one LRU per container, which is how

> > > the CKRM memory controller implemented things about a year ago.

> >

> > Effectively Andrew's idea of faking up nodes is also giving per

> > container LRUs.

>

> Yes, but the NUMA emulation approach is using fixed size containers

> where the size is selectable at the kernel command line, while the CKRM

> (and pzone) approach provides a more dynamic (and complex) solution.

NUMA emulation does not allow guarantee, only limits. It also doesn't
allow over commit (ove commit issue is present in pzone based approach
also).

>

> / magnus

>

>

> -----

> Using Tomcat but need to do more? Need to support web services, security?

> Get stuff done quickly with pre-integrated technology to make your job easier

> Download IBM WebSphere Application Server v.1.0.1 based on Apache Geronimo

> <http://sel.as-us.falkag.net/sel?cmd=lnk&kid=120709&b id=263057&dat=121642>

>

> ckrm-tech mailing list
> <https://lists.sourceforge.net/lists/listinfo/ckrm-tech>
--

Chandra Seetharaman | Be careful what you choose....
- sekharan@us.ibm.com |you may get it.
