
Subject: Re: [ckrm-tech] [PATCH] BC:
resource beancounters (v4) (added user memory)
Posted by [Chandra Seetharaman](#) on Wed, 13 Sep 2006 22:20:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 2006-09-12 at 18:25 -0700, Rohit Seth wrote:

> On Tue, 2006-09-12 at 18:10 -0700, Chandra Seetharaman wrote:

> > On Tue, 2006-09-12 at 17:39 -0700, Rohit Seth wrote:

> > <snip>

> > > yes, it would be there, but is not heavy, IMO.

> > >

> > > I think anything greater than 1% could be a concern for people who are

> > > not very interested in containers but would be forced to live with them.

> >

> > If they are not interested in resource management and/or containers, i

> > do not think they need to pay.

> > >

>

> Think of a single kernel from a vendor that has container support built

> in.

Ok. Understood.

Here are results of some of the benchmarks we have run in the past
(April 2005) with CKRM which showed no/negligible performance impact in
that scenario.

<http://marc.theaimsgroup.com/?l=ckrm-tech&m=111325064322305&w=2>

<http://marc.theaimsgroup.com/?l=ckrm-tech&m=111385973226267&w=2>

<http://marc.theaimsgroup.com/?l=ckrm-tech&m=111291409731929&w=2>

>

<snip>

> > Not at all. If the container they are interested in is guaranteed, I do

> > not see how apps running outside a container would affect them.

> >

>

> Because the kernel (outside the container subsystem) doesn't know of

The core resource subsystem (VM subsystem for memory) would know about
the guarantees and don't cares, and it would handle it appropriately.

> these guarantees...unless you modify the page allocator to have another

> variant of overcommit memory.

>

<snip>

>

> > No, the reclaimer would free up pages associated with the don't care RGs

> > (as the user don't care about the resource made available to them).
> >
>
> And how will the kernel reclaimer know which RGs are don't care?

By looking into the beancounter associated with the container/RG

--

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