
Subject: Re: [ckrm-tech] [PATCH] BC: resource beancounters (v4) (added user memory)

Posted by [Chandra Seetharaman](#) on Wed, 06 Sep 2006 21:47:34 GMT

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

On Wed, 2006-09-06 at 17:06 +0400, Kirill Korotaev wrote:

> Balbir Singh wrote:

> > Kirill Korotaev wrote:

> >

> >> Core Resource Beancounters (BC) + kernel/user memory control.

> >>

> >> BC allows to account and control consumption

> >> of kernel resources used by group of processes.

> >>

> >> Draft UBC description on OpenVZ wiki can be found at

> >> http://wiki.openvz.org/UBC_parameters

> >>

> >> The full BC patch set allows to control:

> >> - kernel memory. All the kernel objects allocatable

> >> on user demand should be accounted and limited

> >> for DoS protection.

> >> E.g. page tables, task structs, vmas etc.

> >

> >

> > One of the key requirements of resource management for us is to be able to

> > migrate tasks across resource groups. Since bean counters do not associate

> > a list of tasks associated with them, I do not see how this can be done

> > with the existing bean counters.

> It was discussed multiple times already.

> The key problem here is the objects which do not belong to tasks.

> e.g. IPC objects. They exist in global namespace and can't be reaccounted.

> At least no one proposed the policy to reaccount.

> And please note, IPCs are not the only such objects.

>From implementation point of view I do not see it to be any different than how it can be done under UBC.

AFAICS, beancounters are associated with tasks not those "objects".

Those "objects" get their bc through some association with a task. The same can be done in the other case also.

If my understanding is wrong, please tell me how one can associate such "object" to a bc.

>

> But I guess your comment mostly concerns user pages, yeah?

> In this case reaccounting can be easily done using page beancounters

> which are introduced in this patch set.

> So if it is a requirement, then lets cooperate and create such functionality.

hmm... that is what I thought I was doing when I was replying on these threads. May be I should have waited for this "call for co-operation" before jumping on it :)

- >
- > So for now I see 2 main requirements from people:
- > - memory reclamation
 - > - tasks moving across beancounters

Please consider the requirements I listed before
<http://marc.theaimsgroup.com/?l=ckrm-tech&m=115593001810616&w=2>

- >
- > I agree with these requirements and lets move into this direction.
- > But moving so far can't be done without accepting:
- > 1. core functionality
 - > 2. accounting

I agree that discussion need to happen on the core functionality and interface.

- >
- > Thanks,
- > Kirill
- >
- >
- > -----
- > 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.comyou may get it.
