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

Posted by [Chandra Seetharaman](#) on Mon, 11 Sep 2006 18:44:22 GMT

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

On Mon, 2006-09-11 at 10:56 +0400, Pavel Emelianov wrote:

<snip>

> >> I think of it as: "I will be allowed to use this many total pages, and
> >> they are guaranteed not to fail." (1), I think. The sum of all of the
> >> system's guarantees must be less than or equal to the amount of free
> >> memory on the machine.

> >

> > Yes, totally agree.

>

> Such a guarantee is really a limit and this limit is even harder than
> BC's one :)

>

> E.g. I have a node with 1Gb of ram and 10 containers with 100Mb
> guarantee each.

In the first place system administrator should not be configuring it that way, Then they are using it as a strict hard limit than guarantee (as the resources guaranteed to one container is not available to others).

Besides, the above configuration is clearly not work conservative.

They should use both guarantee and limit to associate resources to a container/RG.

> I want to start one more. What shall I do not to break guarantees?

CKRM/RG handles it this way:

Amount of a resource a child RG gets is the ratio of its share value to the parent's total # of shares. Children's resource allocation can be changed just by changing the parent's total # of shares.

If you case about initial situation would be:

Total memory in the system 100MB

parent's total # of shares: 100 (1 share == 1MB)

10 children with # of shares: 10 (i.e each children has 10MB)

When I want to add another child, just change parent's total # of shares to be say 125:

Total memory in the system 100MB

parent's total # of shares: 125 (1 share == 0.8MB)
10 children with # of shares: 10 (i.e each children has 8MB)
Now you are left with 25 shares (or 20MB) that you can assign to new
child(ren) as you please.

<snip>

--

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