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

Posted by [Balbir Singh](#) on Tue, 12 Sep 2006 12:01:43 GMT

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

Pavel Emelianov wrote:

<snip>

>>>> E.g. I have a node with 1Gb of ram and 10 containers with 100Mb
>>>> guarantee each.
>>>> I want to start one more. What shall I do not to break guarantees?
>>> Don't start the new container or change the guarantees of the existing
>>> ones
>>> to accommodate this one :) The QoS design (done by the administrator)
>>> should
>>> take care of such use-cases. It would be perfectly ok to have a
>>> container
>>> that does not care about guarantees to set their guarantee to 0 and set
>>> their limit to the desired value. As Chandra has been stating we
>>> need two
>>> parameters (guarantee, limit), either can be optional, but not both.
>> If I set up 9 groups to have 100Mb limit then I have 100Mb assured (on
>> 1Gb node)
>> for the 10th one exactly. And I do not have to set up any guarantee as
>> it won't affect
>> anything. So what a guarantee parameter is needed for?
>> This use case works well for providing guarantee to one container.
>> What if
>> I want guarantees of 100Mb and 200Mb for two containers? How do I setup
>> the system using limits?
> You may set any value from 100 up to 800 Mb for the first one and
> 200-900Mb for
> the second. In case of no other groups first will receive its 100Mb for
> sure and
> so does the second. If there are other groups - their guarantees should
> be concerned.

If I add another group with a guarantee of 100MB, then its limit will be anywhere between 100MB-800MB ?

I do not understand the guarantees being concerned part.

>> Even I restrict everyone else to 700Mb. With this I cannot be sure that
>> the remaining 300Mb will be distributed as 100Mb and 200Mb.
> There's no "everyone else" here - we're talking about a "static" case.
> When new group arrives we need to recalculate guarantees as you said.

I was speaking in general where we have 'n' groups, so thats why I had

"everyone else".

> And here's my next question - what to do if the new guarantee would become
> lower than current amount of unreclaimable memory in BC?
>

I am not quite sure I understand this question. Let me try with an example.
Let's say I have 1GB of memory and I have guarantees of 100 and 200MB for
two groups. Let's say the total unreclaimable memory is 100MB. If I add a new
group with guarantee of 50MB, which is lower than the total amount of
total unreclaimable memory. The addition of the new group should be work
fine since we have $1\text{GB} - 100 - 100 - 200\text{MB}$ of memory available.

--

Balbir Singh,
Linux Technology Center,
IBM Software Labs
