Subject: Re: [ckrm-tech] [PATCH] BC: resource beancounters (v4) (added user memory)
Posted by Pavel Emelianov on Wed, 13 Sep 2006 08:06:41 GMT
View Forum Message <> Reply to Message

Chandra Seetharaman wrote:
> On Tue, 2006-09-12 at 14:48 +0400, Pavel Emelianov wrote:
> <snip>
>
>>> I do not think it is that simple since
>>> - there is typically more than one class I want to set guarantee to
>>> - I will not able to use both limit and guarantee
>>> - Implementation will not be work-conserving.
>>>
>>> Also, How would you configure the following in your model ?
>>>
>>> 5 classes: Class A(10, 40), Class B(20, 100), Class C (30, 100), Class D
>>> (5, 100), Class E(15, 50); (class_name(guarantee, limit))
>>>
>>>
>> What's the total memory amount on the node? Without it it's hard to make
>> any
>> guarantee.
>>
>
> I wrote the example treating them as %, so 100 would be the total amount
> of memory.
>
OK. Then limiting must be done this way (unreclaimable limit/total limit)
A (15/40)
B (25/100)
C (35/100)
D (10/100)
E (20/50)
In this case each group will receive it's guarantee for sure.

E.g. even if A, B, E and D will eat all it's unreclaimable memory then
we'll have
100 - 15 - 25 - 20 - 10 = 30% of memory left (maybe after reclaiming) which
is perfectly enough for C's guarantee.
>
>>> "Limit only" approach works for DoS prevention. But for providing QoS
>>> you would need guarantee.
>>>
>>>
>> You may not provide guarantee on physycal resource for a particular group
>> without limiting its usage by other groups. That's my major idea.
>>

>
> I agree with that, but the other way around (i.e provide guarantee for
> everyone by imposing limits on everyone) is what I am saying is not
> possible.
Then how do you make sure that memory WILL be available when the group needs
it without limiting the others in a proper way?