Subject: Re: [ckrm-tech] [PATCH] BC: resource beancounters (v4) (added user memory)
Posted by Pavel Emelianov on Wed, 13 Sep 2006 13:35:30 GMT
View Forum Message <> Reply to Message

Srivatsa Vaddagiri wrote:
> On Wed, Sep 13, 2006 at 12:06:41PM +0400, Pavel Emelianov wrote:
>> OK. Then limiting must be done this way (unreclaimable limit/total limit)
>> A (15/40)
>> B (25/100)
>> C (35/100)
>
> s/35/30?
Hmmm... No, it must be 35. It IS higher than guarantee you proposed,
but that's OK to have a limit higher than guarantee, isn't it?
>
> Also the different b/n total and unreclaimable limits goes towards
> limiting reclaimable memory i suppose? And 1st limit seems to be a
> hard-limit while the 2nd one is soft?
The first limit (let's call it soft one) is limit for unreclaimable
memory, the second (hard limit) - for booth reclaimable and not.

The ploicy is
1. if BC tries to *mmap()* unreclaimable region (e.g. w/o backed
   file as moving page to swap is not a pure "reclamation") then
   check the soft limit and prohibit mapping in case it is hit;
2. if BC tries to *touch* a page - then check for the hard limit
   and start reclaiming this BC's pages if the limit is hit.

That's how guarantees can be met. Current BC code does perform the
first check and gives you all the levers for the second one - just
the patch(es) with reclamation mechanism is required.
>
>> 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.
>
> I agree by carefully choosing these limits, we can provide some sort of
> QoS, which is a good step to begin with.
Sure. As I've said - soft limiting is already done with BC patches, the
hard one is not prohibited by BC (BCs even prepare a good pad for it).
When reclaiming is done we'll have a hard limit described above.