Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by dev on Fri, 18 Aug 2006 08:47:57 GMT
View Forum Message <> Reply to Message

Rohit Seth wrote:
> On Thu, 2006-08-17 at 17:35 +0400, Kirill Korotaev wrote:
>
>
>>>My preference would be to have container (I keep on saying container,
>>>but resource beancounter) pointer embeded in task, mm(not sure),
>>>address_space and anon_vma structures. This should allow us to track
>>>user land pages optimally. But for tracking kernel usage on behalf of
>>>user, we will have to use an additional field (unless we can re-use
>>>mapping). Please correct me if I'm wrong, though all the kernel
>>>resources will be allocated/freed in context of a user process. And at
>>>that time we know if a allocation should succeed or not. So we may
>>>actually not need to track kernel pages that closely. We are not going
>>>to run reclaim on any of them anyways.
>>
>>objects are really allocated in process context
>>(except for TCP/IP and other softirqs which are done in arbitrary
>>process context!)
>
>
> Can these pages be tagged using mapping field of the page struct.
kernel pages can be taged with mapping field.
User pages - not. So we introduce 2 pointers in the unoin:
union {
  page_ub // for kernel pages
  page_pb // for user pages
}

>
>
>>And objects are not always freed in correct context (!).
>>
>
> You mean beyond Networking and softirq.
>
>
>>Note, page_ub is not for _user_ pages. user pages accounting will be added
>>in next patch set. page_ub is added to track kernel allocations.
>>
>
>
> But will the page_ub be used for some purpose for user land pages?
yes. see above.

Kirill