Subject: Re: [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by dev on Mon, 21 Aug 2006 10:41:28 GMT
View Forum Message <> Reply to Message

>>1. reclaiming user resources is not that good idea as it looks to you.
>>such solutions end up with lots of resources spent on reclaim.
>>for user memory reclaims mean consumption of expensive disk I/O bandwidth
>>which reduces overall system throughput and influences other users.
>>
>
>
> May be I'm overlooking something very obvious. Please tell me, what
> happens when a user hits a page fault and the page allocator is easily
> able to give a page from its pcp list. But container is over its limit
> of physical memory. In your patch there is no attempt by container
> support to see if some of the user pages are easily reclaimable. What
> options a user will have to make sure some room is created.
The patch set send doesn't control user memory!
This topic is about kernel memory...

>>2. kernel memory is mostly not reclaimable. can you reclaim vma structs or ipc ids?
>
>
> I'm not arguing about that at all. If people want to talk about
> reclaiming kernel pages then that should be done independent of this
> subject.
Then why do you mess user pages accounting into this thread then?

Kirill