Subject: Re: [RFC][PATCH] another swap controller for cgroup
Posted by Daisuke Nishimura on Tue, 25 Mar 2008 04:35:53 GMT
View Forum Message <> Reply to Message

YAMAMOTO Takashi wrote:
> hi,
>
>> Daisuke Nishimura wrote:
>>> Hi, Yamamoto-san.
>>>
>>> I'm reviewing and testing your patch now.
>>>
>> In building kernel infinitely(in a cgroup of
>> memory.limit=64M and swap.limit=128M, with swappiness=100),
>> almost all of the swap (1GB) is consumed as swap cache
>> after a day or so.
>> As a result, processes are occasionally OOM-killed even when
>> the swap.usage of the group doesn't exceed the limit.
>>
>> I don't know why the swap cache uses up swap space.
>> I will test whether a similar issue happens without your patch.
>> Do you have any thoughts?
>
> my patch tends to yield more swap cache because it makes try_to_unmap
> fail and shrink_page_list leaves swap cache in that case.
> i'm not sure how it causes 1GB swap cache, tho.
>

Agree.

I suspected that the cause of this problem was the behavior
of shrink_page_list as you said, so I thought one of Rik's
split-lru patchset:

  http://lkml.org/lkml/2008/3/4/492
  [patch 04/20] free swap space on swap-in/activation

would reduce the usage of swap cache to half of the total swap.
But it didn't help, so I think there may be some other causes.

Thanks,
Daisuke Nishimura.

_____