Subject: Re: [RFC][PATCH] another swap controller for cgroup
Posted by yamamoto on Tue, 25 Mar 2008 08:57:23 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.

do you mean you tested with the patch in the url?
i don't think remove_exclusive_swap_page works for us
because our page has more references than it expects.
ie. ptes, cache, isolate_page

(unless your tree has another change for remove_exclusive_swap_page.
i haven't checked other patches in the patchset.)

YAMAMOTO Takashi

>
>
> Thanks,
> Daisuke Nishimura.

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers