Subject: Re: [RFD][PATCH] memcg: Move Usage at Task Move
Posted by KAMEZAWA Hiroyuki on Wed, 11 Jun 2008 04:05:32 GMT
View Forum Message <> Reply to Message

On Wed, 11 Jun 2008 12:45:14 +0900 (JST)
yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:

> > > having said that, if you decide to put too large tasks into
> > > a cgroup with too small limit, i don't think that there are
> > > many choices besides OOM-kill and allowing "exceed".
> > >
> > IMHO, allowing exceed is harmfull without changing the definition of "limit".
> > "limit" is hard-limit, now, not soft-limit. Changing the defintion just for
> > this is not acceptable for me.
>
> even with the current code, the "exceed" condition can be created
> by simply lowering the limit.
> (well, i know that some of your patches floating around change it.)
>
Yes, I write it now ;) Handling exceed contains some troubles

 - when resizing limit, to what extent exceed is allowed ?
 - Once exceed, no new page allocation can success and
   some random process will die because of OOM.

> > Maybe "move" under limit itself is crazy ops....Hmm...
> >
> > Should we allow task move when the destination cgroup is unlimited ?
> > Isn't it useful ?
>
> i think it makes some sense.
>
> > > actually, i think that #3 and #5 are somewhat similar.
> > > a big difference is that, while #5 shrinks the cgroup immediately,
> > > #3 does it later. in case we need to do OOM-kill, i prefer to do it
> > > sooner than later.
> > >
> > #3 will not cause OOM-killer, I hope...A user can notice memory shortage.
>
> we are talking about the case where a cgroup's working set is getting
> hopelessly larger than its limit. i don't see why #3 will not
> cause OOM-kill. can you explain?
>
just because #3 doesn't move resource, just drop.

Thanks,
-Kame

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