

---

Subject: Re: [PATCH 0/4] swapcgroup(v2)  
Posted by [Rik van Riel](#) on Fri, 23 May 2008 03:32:07 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Fri, 23 May 2008 12:10:27 +0900  
KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> wrote:  
> On Thu, 22 May 2008 22:26:55 -0400  
> Rik van Riel <riel@redhat.com> wrote:  
>  
> > Even worse is that a cgroup has NO CONTROL over how much  
> > of its memory is kept in RAM and how much is swapped out.  
> > Could you explain "NO CONTROL" ? cgroup has LRU....  
> > 'how much memory should be swapped out from memory' is well controlled  
> > in the VM besides LRU logic ?

The kernel controls what is swapped out. The userland processes in the cgroup can do nothing to reduce their swap usage.

> Consider following system. (and there is no swap controller.)  
> Memory 4G. Swap 1G. with 2 cgroups A, B.  
>  
> state 1) swap is not used.  
> A....memory limit to be 1G no swap usage memory\_usage=0M  
> B....memory limit to be 1G no swap usage memory\_usage=0M  
>  
> state 2) Run a big program on A.  
> A....memory limit to be 1G and try to use 1.7G. uses 700MBytes of swap.  
> memory\_usage=1G swap\_usage=700M  
> B....memory\_usage=0M  
>  
> state 3) A some of programs ends in 'A'  
> A....memory\_usage=500M swap\_usage=700M  
> B....memory\_usage=0M.  
>  
> state 4) Run a big program on B.  
> A...memory\_usage=500M swap\_usage=700M.  
> B...memory\_usage=1G swap\_usage=300M  
>  
> Group B can only use 1.3G because of unfair swap use of group A.  
> But users think why A uses 700M of swap with 500M of free memory....  
>  
> If we don't have limitation to swap, we'll have to innovate a way to move swap  
> to memory in some reasonable logic.

OK, I see the use case.

In the above example, it would be possible for cgroup A

to have only 800MB of anonymous memory total, in addition to 400MB of page cache. The page cache could push the anonymous memory into swap, indirectly penalizing how much memory cgroup B can use.

Of course, it could be argued that the system should just be run with enough swap space, but that is another story :)

--

All rights reversed.

---

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

---