

---

Subject: Re: A consideration on memory controller.  
Posted by [KAMEZAWA Hiroyuki](#) on Mon, 21 Jan 2008 10:22:23 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Mon, 21 Jan 2008 15:20:36 +0530  
Balbir Singh <[balbir@linux.vnet.ibm.com](mailto:balbir@linux.vnet.ibm.com)> wrote:

> > > This topic needs more discussion, we have some form of page-cache  
> > > control built into the memory controller.  
> > >  
> > Hmm. ok. I'm looking forward to see.  
> >  
>  
> Could you elaborate on what sort of page-cache control you need, is it  
> global page-cache control?  
>

What I mentioned to in my mail was global page cache.

But, my purpose is isolate workload between applications and guarantee  
stable performance/latency and memory controller's limiting page usage  
feature will do half of work.

So, per-cgroup page-cache control is (maybe) good enough for making applications  
stable by keeping room for immediate use of anon and by avoiding unnecessary swapout.

but above can be achieved by

- reserve/guarantee amount of free memory by some technique, background  
kthread, throttling,
- do good design of swappiness.
- etc...

So, the word "limiting page cache" itself is not important. Above will  
limit amount of page-cache as a result.

Ways to reach my goal is not one, maybe. I should find the best one :)

Thanks,  
-Kame

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---