Subject: Re: [RFC] memory controller : backgorund reclaim and avoid excessive locking [5/5] lazy page_cgroup
Posted by KAMEZAWA Hiroyuki on Mon, 18 Feb 2008 04:54:03 GMT

On Mon, 18 Feb 2008 10:05:16 +0530
Balbir Singh <balbir@linux.vnet.ibm.com> wrote:
> Hi, KAMEZAWA-San,
>
> For this and the next patch, do you know if there is a performance improvement?
> Have we measured it. The complexity of the code seems to be high, want to make
> sure it is worth it.

I'd like to make codes simpler (by removing knobs and...)
I confirmed lock-contention/cache bouncing was much reduced.
(you can see it by /proc/lock_stat)
But I have no number now. I'm now searching an apps which can show performance
improvement clearly and wait for avilable machine ;)
So, please wait. I'm not hurrying.

> Could we re-use the pagevec mechansim for LRU handling?
Hmm, it's also complicated.

>That was my long term plan and I also wanted to move the kmalloc's to their
>own cache and use kmem_cache_alloc and try some other experiments.
maybe good idea.

> May be batching the alloc's for
> page_container can be done in the same way that you've proposed LRU and
> accounting changes.
>
Maybe.

Thanks,
-Kame

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