
Subject: Re: [RFC][PATCH][0/4] Memory controller (RSS Control)
Posted by [Balbir Singh](#) on Mon, 19 Feb 2007 10:39:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

Paul Menage wrote:

> On 2/19/07, Andrew Morton <akpm@linux-foundation.org> wrote:

>>

>> Alas, I fear this might have quite bad worst-case behaviour. One small
>> container which is under constant memory pressure will churn the
>> system-wide LRUs like mad, and will consume rather a lot of system time.
>> So it's a point at which container A can deleteriously affect things

>> which

>> are running in other containers, which is exactly what we're supposed to
>> not do.

>

> I think it's OK for a container to consume lots of system time during
> reclaim, as long as we can account that time to the container involved
> (i.e. if it's done during direct reclaim rather than by something like
> kswapd).

>

> Churning the LRU could well be bad though, I agree.

>

I completely agree with you on reclaim consuming time.

Churning the LRU can be avoided by the means I mentioned before

1. Add a container pointer (per page struct), it is also useful for the page cache controller
2. Check if the page belongs to a particular container before the `list_del(&page->lru)`, so that those pages can be skipped.
3. Use a double LRU list by overloading the `lru list_head` of struct page.

> Paul

>

--

Warm Regards,
Balbir Singh
