
Subject: Re: [ckrm-tech] [RFC][PATCH][0/4] Memory controller (RSS Control)
Posted by [dev](#) on Mon, 19 Feb 2007 09:46:15 GMT

[View Forum Message](#) <> [Reply to Message](#)

> 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).

hmm, is it ok to scan 100Gb of RAM for 10MB RAM container?

in UBC patch set we used page beanccounters to track container pages.

This allows to make efficient scan contoler and reclamation.

Thanks,
Kirill
