
Subject: Re: [RFC][PATCH][0/4] Memory controller (RSS Control)
Posted by [Magnus Damm](#) on Mon, 19 Feb 2007 09:16:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 2/19/07, Andrew Morton <akpm@linux-foundation.org> wrote:
> On Mon, 19 Feb 2007 12:20:19 +0530 Balbir Singh <balbir@in.ibm.com> wrote:
>
> > This patch applies on top of Paul Menage's container patches (V7) posted at
> >
> > <http://lkml.org/lkml/2007/2/12/88>
> >
> > It implements a controller within the containers framework for limiting
> > memory usage (RSS usage).

> The key part of this patchset is the reclaim algorithm:
>
> 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.

Nice with a simple memory controller. The downside seems to be that it doesn't scale very well when it comes to reclaim, but maybe that just comes with being simple. Step by step, and maybe this is a good first step?

Ideally I'd like to see unmapped pages handled on a per-container LRU with a fallback to the system-wide LRUs. Shared/mapped pages could be handled using PTE ageing/unmapping instead of page ageing, but that may consume too much resources to be practical.

/ magnus
