
Subject: Re: [RFC][PATCH][0/4] Memory controller (RSS Control)
Posted by [Andrew Morton](#) on Mon, 19 Feb 2007 08:54:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

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).

It's good to see someone building on someone else's work for once, rather than everyone going off in different directions. It makes one hope that we might actually achieve something at last.

The key part of this patchset is the reclaim algorithm:

```
> @@ -636,6 +642,15 @@ static unsigned long isolate_lru_pages(u
>
> list_del(&page->lru);
> target = src;
> + /*
> +  * For containers, do not scan the page unless it
> +  * belongs to the container we are reclaiming for
> +  */
> + if (container && !page_in_container(page, zone, container)) {
> + scan--;
> + goto done;
> + }
```

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.
