
Subject: Re: [RFC][PATCH 2/7] RSS controller core
Posted by [Herbert Poetzl](#) on Fri, 09 Mar 2007 16:48:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Mar 06, 2007 at 02:00:36PM -0800, Andrew Morton wrote:

> On Tue, 06 Mar 2007 17:55:29 +0300

> Pavel Emelianov <xemul@sw.ru> wrote:

>

> > +struct rss_container {

> > + struct res_counter res;

> > + struct list_head page_list;

> > + struct container_subsys_state css;

> > +};

> > +

> > +struct page_container {

> > + struct page *page;

> > + struct rss_container *cnt;

> > + struct list_head list;

> > +};

>

> ah. This looks good. I'll find a hunk of time to go through this work
> and through Paul's patches. It'd be good to get both patchsets lined
> up in -mm within a couple of weeks. But..

doesn't look so good for me, mainly becaus of the
additional per page data and per page processing

on 4GB memory, with 100 guests, 50% shared for each
guest, this basically means ~1mio pages, 500k shared
and 1500k x sizeof(page_container) entries, which
roughly boils down to ~25MB of wasted memory ...

increase the amount of shared pages and it starts
getting worse, but maybe I'm missing something here

> We need to decide whether we want to do per-container memory
> limitation via these data structures, or whether we do it via a
> physical scan of some software zone, possibly based on Mel's patches.

why not do simple page accounting (as done currently
in Linux) and use that for the limits, without
keeping the reference from container to page?

best,
Herbert

>

> Containers mailing list

> Containers@lists.osdl.org
> <https://lists.osdl.org/mailman/listinfo/containers>

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
