
Subject: Re: [RFC][PATCH 2/7] RSS controller core

Posted by [dev](#) on Sun, 11 Mar 2007 12:13:08 GMT

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

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..

>

> 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.

i.e. a separate memzone for each container?

imho memzone approach is inconvenient for pages sharing and shares accounting.

it also makes memory management more strict, forbids overcommitting
per-container etc.

Maybe you have some ideas how we can decide on this?

Thanks,

Kirill
