Subject: Re: [PATCH 8/8] Per-container pages reclamation Posted by Balbir Singh on Thu, 24 May 2007 07:59:55 GMT

View Forum Message <> Reply to Message

```
Pavel Emelianov wrote:
>> Index: linux-2.6.20/mm/rss_container.c
>> --- linux-2.6.20.orig/mm/rss_container.c 2007-05-15 05:13:46.000000000 -0700
>> +++ linux-2.6.20/mm/rss container.c 2007-05-16 20:45:45.000000000 -0700
>> @ @ -212,6 +212,7 @ @ void container rss del(struct page conta
>>
>> css_put(&rss->css);
>> kfree(pc);
>> + init_page_container(page);
> This hunk is bad.
> See, when the page drops its mapcount to 0 it may be reused right
> after this if it belongs to a file map - another CPU can touch it.
> Thus you're risking to reset the wrong container.
>
> The main idea if the accounting is that you cannot trust the
> page container(page) value after the page's mapcount became 0.
>
Good catch, I'll move the initialization to free_hot_cold_page().
I'm attaching a new patch. I've also gotten rid of the unused
variable page in container_rss_del().
I've compile and boot tested the fix
```

Thanks, Balbir Singh Linux Technology Center IBM, ISTL