
Subject: Re: [RFC][PATCH 4/7] RSS accounting hooks over the code
Posted by [ebiederm](#) on Tue, 13 Mar 2007 09:43:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dave Hansen <hansenc@us.ibm.com> writes:

> On Mon, 2007-03-12 at 20:07 +0300, Kirill Korotaev wrote:
>> > On Mon, 2007-03-12 at 19:23 +0300, Kirill Korotaev wrote:
>> >> For these you essentially need per-container page->_mapcount counter,
>> >> otherwise you can't detect whether rss group still has the page in question
> being mapped
>> >> in its processes' address spaces or not.
>> >
>> > What do you mean by this? You can always tell whether a process has a
>> > particular page mapped. Could you explain the issue a bit more. I'm
>> > not sure I get it.
>> When we do charge/uncharge we have to answer on another question:
>> "whether *any* task from the *container* has this page mapped", not the
>> "whether *this* task has this page mapped".
>
> That's a bit more clear. ;)
>
> OK, just so I make sure I'm getting your argument here. It would be too
> expensive to go looking through all of the rmap data for _any_ other
> task that might be sharing the charge (in the same container) with the
> current task that is doing the unmapping.

Which is a questionable assumption. Worse case we are talking a list
several thousand entries long, and generally if you are used by the same
container you will hit one of your processes long before you traverse
the whole list.

So at least the average case performance should be good.

It is only in the case when you a page is shared between multiple
containers when this matters.

Eric

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