
Subject: Re: [RFC][PATCH 4/7] RSS accounting hooks over the code
Posted by [Ethan Solomita](#) on Wed, 28 Mar 2007 20:15:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Nick Piggin wrote:

> Eric W. Biederman wrote:

>> First touch page ownership does not guarantee give me anything useful
>> for knowing if I can run my application or not. Because of page
>> sharing my application might run inside the rss limit only because
>> I got lucky and happened to share a lot of pages with another running
>> application. If the next I run and it isn't running my application
>> will fail. That is ridiculous.

>

> Let's be practical here, what you're asking is basically impossible.

>

> Unless by deterministic you mean that it never enters the a non
> trivial syscall, in which case, you just want to know about maximum
> RSS of the process, which we already account).

If we used Beancounters as Pavel and Kirill mentioned, that would keep track of each container that has referenced a page, not just the first container. It sounds like beancounters can return a usage count where each page is divided by the number of referencing containers (e.g. 1/3rd if 3 containers share a page). Presumably it could also return a full count of 1 to each container.

If we look at data in the latter form, i.e. each container must pay fully for each page used, then Eric could use that to determine real usage needs of the container. However we could also use the fractional count in order to do things such as charging the container for its actual usage. i.e. full count for setting guarantees, fractional for actual usage.

-- Ethan

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
