
Subject: Re: [PATCH 5/8] RSS accounting hooks over the code

Posted by [Balbir Singh](#) on Fri, 01 Jun 2007 06:48:10 GMT

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

Andrew Morton wrote:

> On Wed, 30 May 2007 19:34:18 +0400

> Pavel Emelianov <xemul@openvz.org> wrote:

>

>> As described above, pages are charged to their first touchers.

>> The first toucher is determined using pages' _mapcount

>> manipulations in rmap calls.

>>

>> Page is charged in two stages:

>> 1. preparation, in which the resource availability is checked.

>> This stage may lead to page reclamation, thus it is performed

>> in a "might-sleep" places;

>> 2. the container assignment to page. This is done in an atomic

>> code that handles races between multiple touchers.

>

> I suppose we need to think about what to do about higher-order pages, and

> compound pages, and hugetlb memory.

Yes, that needs to be done eventually. I feel that support can be added incrementally once we have a stable controller.

--

Warm Regards,

Balbir Singh

Linux Technology Center

IBM, ISTL
