

---

Subject: Re: [PATCH 8/8] Per-container pages reclamation  
Posted by [Pavel Emelianov](#) on Fri, 01 Jun 2007 07:40:12 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Andrew Morton wrote:

> On Thu, 31 May 2007 14:35:43 +0400

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

>

>>> Would I be correct in guessing that pages which are on the  
>>> per-rss-container lists are also eligible for reclaim off the traditional  
>>> page LRUs? If so, how does that work? When a page gets freed off the  
>> Yes. All the pages are accessible from booth - global and per-container  
>> LRU lists and reclamation can be performed from booth.

>>

>>> per-zone LRUs does it also get removed from the per-rss\_container LRU? But  
>>> how can this be right?

>> I don't get your idea here.

>

> If we have a page which is on the zone LRU with refcount=1 and someone does  
> put\_page() on it, we will take that page off the zone LRU and then actually  
> free the page.

>

> I am assuming that your patches change that logic so that we will also  
> remove that page from the per-container LRU at the same time?

Oh, I see. No that will work another way. Page stays in per-container LRU lists as long as it is mapped. When the last process is unmapping the page, it is removed, but stays in global LRU till its refcount drops to 0.

Since all unmaps happen from process-context this looks to be safe.

> If so, the problem which I see is that, under rare circumstances, that  
> final put\_page() will occur from interrupt context. Hence we would be  
> trying to remove the page from the per-container LRU at interrupt time.  
> But the locks which you have in there are not suited for taking from  
> interrupt context.

>

>

Thanks,  
Pavel.

---