
Subject: Re: [-mm PATCH 0/7] Memory controller introduction
Posted by [Balbir Singh](#) on Thu, 05 Jul 2007 14:32:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi, Pavel,

Pavel Emelianov wrote:

> Balbir Singh wrote:

>

> As far as I remember at OLS we decided to implement per-zone RLU
> lists and reuse the lru lock as well. This will remove all the
> problems with per-container lists inconsistency.

>

It's there in the TODO list. It is easy to implement. I can do that
in the next revision. We are re-using the LRU lock, since the
isolate_pages callback is called under the zone's LRU lock.

> Separate limits for RSS and RSS+pagecache are also a must.

>

I remember that we discussed having one limit, but we can come up
with a configuration parameter to do it. I'll do that in the
next release.

> BTW, if you send smb. else's patches you may include a 'From: xxx'
> line into the letter to address the original author.

>

I'll do that for the res_counters_infra patch. Although I had used
res_counter, the hooks in mm_struct earlier (<http://lkml.org/lkml/2007/2/19/12>).

A lot of the code for the reclaim logic and the meta_page is yours.
Please do identify patches where you would like to see your
signed-off-by and where you think the entire patch is completely
yours.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

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