
Subject: Re: [RFC] [-mm PATCH] Memory controller fix swap charging context in
unuse_pte()

Posted by [Balbir Singh](#) on Mon, 29 Oct 2007 22:01:38 GMT

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

Hugh Dickins wrote:

> On Mon, 29 Oct 2007, Balbir Singh wrote:

>> On Mon, Oct 29, 2007 at 01:57:40AM +0530, Balbir Singh wrote:

>> Hugh Dickins wrote:

>>

>> [snip]

>>

>>> Without your mem_cgroup mods in mm/swap_state.c, unuse_pte makes
>>> the right assignments (I believe). But I find that swapout (using
>>> 600M in a 512M machine) from a 200M cgroup quickly OOMs, whereas
>>> it behaves correctly with your mm/swap_state.c.

>>>

>> On my UML setup, I booted the UML instance with 512M of memory and
>> used the swapout program that you shared. I tried two things

>>

>>

>> 1. Ran swapout without any changes. The program ran well without

>> any OOM condition occuring, lot of reclaim occured.

>> 2. Ran swapout with the changes to mm/swap_state.c removed (diff below)

>> and I still did not see any OOM. The reclaim count was much lesser

>> since swap cache did not get accounted back to the cgroup from

>> which pages were being evicted.

>>

>> I am not sure why I don't see the OOM that you see, still trying. May be

>> I missing something obvious at this late hour in the night :-)

>

> I reconfirm that I do see those OOMs. I'll have to try harder to

> analyze how they come about: I sure don't expect you to debug a

> problem you cannot reproduce. But what happens if you try it

> native rather than using UML?

>

> Hugh

On a real box - a powerpc machine that I have access to

1. I don't see the OOM with the mods removed (I have swap space at-least twice of RAM - with mem=512M, I have at-least 1G of swap).
2. Running under the container is much much faster than running swapout in the root container. The machine is almost unusable if swapout is run under the root container

At this momemnt, I suspect one of two things

1. Our mods to swap_state.c are different
2. Our configuration is different, main-memory to swap-size ratio

--

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>
