
Subject: Re: [RFC][-mm PATCH 6/8] Memory controller add per container LRU and reclaim (v3)

Posted by [Balbir Singh](#) on Tue, 24 Jul 2007 12:14:09 GMT

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

YAMAMOTO Takashi wrote:

```
> hi,
>
>> +unsigned long mem_container_isolate_pages(unsigned long nr_to_scan,
>> +    struct list_head *dst,
>> +    unsigned long *scanned, int order,
>> +    int mode, struct zone *z,
>> +    struct mem_container *mem_cont,
>> +    int active)
>> +{
>> + unsigned long nr_taken = 0;
>> + struct page *page;
>> + unsigned long scan;
>> + LIST_HEAD(mp_list);
>> + struct list_head *src;
>> + struct meta_page *mp;
>> +
>> + if (active)
>> +   src = &mem_cont->active_list;
>> + else
>> +   src = &mem_cont->inactive_list;
>> +
>> + for (scan = 0; scan < nr_to_scan && !list_empty(src); scan++) {
>> +   mp = list_entry(src->prev, struct meta_page, lru);
>> +   page = mp->page;
>> +
>
> - is it safe to pick the lists without mem_cont->lru_lock held?
>
> - what prevents mem_container_uncharge from freeing this meta_page
>   behind us?
>
> YAMAMOTO Takashi
```

Hi, YAMAMOTO,

We do take the lru_lock before deleting the page from the list and in mem_container_move_lists(). But, I guess like you point out page = mp->page might not be a safe operation. I'll fix the problem in the next release.

Thanks for the review,

--

Balbir Singh
Linux Technology Center
IBM, ISTL

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