
Subject: Re: [PATCH 4/4] swapcgroup: modify vm_swap_full for cgroup
Posted by [Daisuke Nishimura](#) on Thu, 22 May 2008 12:34:50 GMT

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

Hi,

On 2008/05/22 15:45 +0900, YAMAMOTO Takashi wrote:

```
>> @@ -1892,3 +1892,36 @@ int valid_swaphandles(swp_entry_t entry, unsigned long *offset)
>> *offset = ++toff;
>> return nr_pages? ++nr_pages: 0;
>> }
>> +
>> +#ifdef CONFIG_CGROUP_SWAP_RES_CTLR
>> +int swap_cgroup_vm_swap_full(struct page *page)
>> +{
>> + int ret;
>> + struct swap_info_struct *p;
>> + struct mem_cgroup *mem;
>> + u64 usage;
>> + u64 limit;
>> + swp_entry_t entry;
>> +
>> + VM_BUG_ON(!PageLocked(page));
>> + VM_BUG_ON(!PageSwapCache(page));
>> +
>> + ret = 0;
>> + entry.val = page_private(page);
>> + p = swap_info_get(entry);
>> + if (!p)
>> + goto out;
>> +
>> + mem = p->memcg[swp_offset(entry)];
>> + usage = swap_cgroup_read_usage(mem) / PAGE_SIZE;
>> + limit = swap_cgroup_read_limit(mem) / PAGE_SIZE;
>> + limit = (limit < total_swap_pages) ? limit : total_swap_pages;
>> +
>> + ret = usage * 2 > limit;
>> +
>> + spin_unlock(&swap_lock);
>> +
>> +out:
>> + return ret;
>> +}
>> +#endif
>
> shouldn't it check the global usage (nr_swap_pages) as well?
>
> YAMAMOTO Takashi
```

>

I didn't check global usage because I didn't want some group to be influenced by other groups.

But in above code, there would be some cases that `vm_swap_full()` returns false even when more than half of swap is used in global.

Thanks you for pointing it out.

How about something like this?

```
:  
usage = swap_cgroup_read_usage(mem); //no need to align to number of page  
limit = swap_cgroup_read_limit(mem); //no need to align to number of page  
ret = (usage * 2 > limit) || (nr_swap_pages * 2 < total_swap_pages)  
:
```

Thanks,
Daisuke Nishimura.

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