
Subject: Re: [-mm PATCH 4/8] Memory controller memory accounting (v2)

Posted by [Balbir Singh](#) on Tue, 10 Jul 2007 08:41:18 GMT

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

On 7/10/07, YAMAMOTO Takashi <yamamoto@valinux.co.jp> wrote:

```
> hi,  
>  
> > diff -puN mm/memory.c~mem-control-accounting mm/memory.c  
> > --- linux-2.6.22-rc6/mm/memory.c~mem-control-accounting 2007-07-05  
13:45:18.000000000 -0700  
> > +++ linux-2.6.22-rc6-balbir/mm/memory.c 2007-07-05 13:45:18.000000000 -0700  
>  
> > @@ -1731,6 +1736,9 @@ gotten:  
> >         cow_user_page(new_page, old_page, address, vma);  
> >     }  
> >  
> > +    if (mem_container_charge(new_page, mm))  
> > +        goto oom;  
> > +  
> >     /*  
> >     * Re-check the pte - we dropped the lock  
> >     */  
>  
> it seems that the page will be leaked on error.
```

You mean meta_page right?

```
>  
> > @@ -2188,6 +2196,11 @@ static int do_swap_page(struct mm_struct  
> >     }  
> >  
> >     delayacct_clear_flag(DELAYACCT_PF_SWAPIN);  
> > +    if (mem_container_charge(page, mm)) {  
> > +        ret = VM_FAULT_OOM;  
> > +        goto out;  
> > +    }  
> > +  
> >     mark_page_accessed(page);  
> >     lock_page(page);  
> >  
>  
> ditto.  
>  
> > @@ -2264,6 +2278,9 @@ static int do_anonymous_page(struct mm_s  
> >         if (!page)  
> >             goto oom;  
> >  
> > +        if (mem_container_charge(page, mm))
```

```
> > +         goto oom;  
> > +  
> >     entry = mk_pte(page, vma->vm_page_prot);  
> >     entry = maybe_mkwrite(pte_mkdirty(entry), vma);  
> >  
>  
> ditto.  
>  
> can you check the rest of the patch by yourself? thanks.  
>
```

Excellent catch! I'll review the accounting framework and post the updated version soon

Balbir

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