
Subject: [PATCH] fix spurious EBUSY on memory cgroup removal

Posted by [yamamoto](#) on Tue, 25 Mar 2008 05:47:13 GMT

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

[resending with To: akpm. Andrew, can you include this in -mm tree?]

hi,

the following patch is to fix spurious EBUSY on cgroup removal.

YAMAMOTO Takashi

call mm_free_cgroup earlier.

otherwise a reference due to lazy mm switching can prevent cgroup removal.

Signed-off-by: YAMAMOTO Takashi <yamamoto@valinux.co.jp>

Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>

--- linux-2.6.24-rc8-mm1/kernel/fork.c.BACKUP 2008-01-23 14:43:29.000000000 +0900

+++ linux-2.6.24-rc8-mm1/kernel/fork.c 2008-01-31 17:26:31.000000000 +0900

@@ -393,7 +393,6 @@ void __mmdrop(struct mm_struct *mm)

```
{
    BUG_ON(mm == &init_mm);
    mm_free_pgd(mm);
- mm_free_cgroup(mm);
    destroy_context(mm);
    free_mm(mm);
}
```

@@ -415,6 +414,7 @@ void mmput(struct mm_struct *mm)

```
    spin_unlock(&mmlist_lock);
}
```

```
    put_swap_token(mm);
+ mm_free_cgroup(mm);
    mmdrop(mm);
}
```

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [PATCH] fix spurious EBUSY on memory cgroup removal

Posted by [akpm](#) on Tue, 25 Mar 2008 05:53:09 GMT

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

On Tue, 25 Mar 2008 14:47:13 +0900 (JST) yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:

> [resending with To: akpm. Andrew, can you include this in -mm tree?]

Shouldn't it be in 2.6.25?

```
> hi,
>
> the following patch is to fix spurious EBUSY on cgroup removal.
>
> YAMAMOTO Takashi
>
>
> call mm_free_cgroup earlier.
> otherwise a reference due to lazy mm switching can prevent cgroup removal.
>
> Signed-off-by: YAMAMOTO Takashi <yamamoto@valinux.co.jp>
> Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>
> ---
>
> --- linux-2.6.24-rc8-mm1/kernel/fork.c.BACKUP 2008-01-23 14:43:29.000000000 +0900
> +++ linux-2.6.24-rc8-mm1/kernel/fork.c 2008-01-31 17:26:31.000000000 +0900
> @@ -393,7 +393,6 @@ void __mmdrop(struct mm_struct *mm)
> {
>     BUG_ON(mm == &init_mm);
>     mm_free_pgdpd(mm);
>     mm_free_cgroup(mm);
>     destroy_context(mm);
>     free_mm(mm);
> }
> @@ -415,6 +414,7 @@ void mmput(struct mm_struct *mm)
>     spin_unlock(&mm_list_lock);
> }
>     put_swap_token(mm);
> + mm_free_cgroup(mm);
>     mmdrop(mm);
> }
> }
```

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

Subject: Re: [PATCH] fix spurious EBUSY on memory cgroup removal
Posted by [Balbir Singh](#) on Tue, 25 Mar 2008 06:28:16 GMT

KAMEZAWA Hiroyuki wrote:

> On Mon, 24 Mar 2008 22:53:09 -0700

> Andrew Morton <akpm@linux-foundation.org> wrote:

>

>> On Tue, 25 Mar 2008 14:47:13 +0900 (JST) yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:

>>

>>> [resending with To: akpm. Andrew, can you include this in -mm tree?]

>> Shouldn't it be in 2.6.25?

>>

> I think this should be.

>

> Thanks,

> -Kame

Me too

--

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>

Subject: Re: [PATCH] fix spurious EBUSY on memory cgroup removal

Posted by [yamamoto](#) on Tue, 25 Mar 2008 06:28:53 GMT

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

> On Tue, 25 Mar 2008 14:47:13 +0900 (JST) yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:

>

> > [resending with To: akpm. Andrew, can you include this in -mm tree?]

>

> Shouldn't it be in 2.6.25?

yes, probably.

(i'm not sure about linux development model.)

YAMAMOTO Takashi

>

> > hi,
> >
> > the following patch is to fix spurious EBUSY on cgroup removal.
> >
> > YAMAMOTO Takashi
> >
> >
> > call mm_free_cgroup earlier.
> > otherwise a reference due to lazy mm switching can prevent cgroup removal.
> >
> > Signed-off-by: YAMAMOTO Takashi <yamamoto@valinux.co.jp>
> > Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>
> > ---
> >
> > --- linux-2.6.24-rc8-mm1/kernel/fork.c.BACKUP 2008-01-23 14:43:29.000000000 +0900
> > +++ linux-2.6.24-rc8-mm1/kernel/fork.c 2008-01-31 17:26:31.000000000 +0900
> > @@ -393,7 +393,6 @@ void __mmdrop(struct mm_struct *mm)
> > {
> > BUG_ON(mm == &init_mm);
> > mm_free_pgd(mm);
> > - mm_free_cgroup(mm);
> > destroy_context(mm);
> > free_mm(mm);
> > }
> > @@ -415,6 +414,7 @@ void mmput(struct mm_struct *mm)
> > spin_unlock(&mmlist_lock);
> > }
> > put_swap_token(mm);
> > + mm_free_cgroup(mm);
> > mmdrop(mm);
> > }
> > }
> >
> _____
> Containers mailing list
> Containers@lists.linux-foundation.org
> <https://lists.linux-foundation.org/mailman/listinfo/containers>

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

Subject: Re: [PATCH] fix spurious EBUSY on memory cgroup removal
Posted by [KAMEZAWA Hiroyuki](#) on Tue, 25 Mar 2008 06:30:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 24 Mar 2008 22:53:09 -0700
Andrew Morton <akpm@linux-foundation.org> wrote:

> On Tue, 25 Mar 2008 14:47:13 +0900 (JST) yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:

>

> > [resending with To: akpm. Andrew, can you include this in -mm tree?]

>

> Shouldn't it be in 2.6.25?

>

I think this should be.

Thanks,

-Kame

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
