
Subject: Re: [-mm PATCH 6/9] Memory controller add per container LRU and reclaim (v4)

Posted by [Peter Zijlstra](#) on Mon, 30 Jul 2007 13:59:20 GMT

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

On Mon, 2007-07-30 at 19:07 +0530, Dhaval Giani wrote:

```
> Hi Balbir,  
>  
> > diff -puN mm/memcontrol.c~mem-control-lru-and-reclaim mm/memcontrol.c  
> > --- linux-2.6.23-rc1-mm1/mm/memcontrol.c~mem-control-lru-and-reclaim 2007-07-28  
01:12:50.000000000 +0530  
> > +++ linux-2.6.23-rc1-mm1-balbir/mm/memcontrol.c 2007-07-28 01:12:50.000000000 +0530  
>  
> > /*  
> > * The memory controller data structure. The memory controller controls both  
> > @@ -51,6 +54,10 @@ struct mem_container {  
> > */  
> > struct list_head active_list;  
> > struct list_head inactive_list;  
> > +/*  
> > + * spin_lock to protect the per container LRU  
> > + */  
> > + spinlock_t lru_lock;  
> > };  
>  
> The spinlock is not annotated by lockdep. The following patch should do  
> it.
```

One does not need explicit lockdep annotations unless there is a non obvious use of the locks. A typical example of that would be the inode locks, that get placed differently in the various filesystem's locking hierarchy and might hence seem to generate contradictory locking rules - even though they are consistent within a particular filesystem.

So unless there are 2 or more distinct locking hierarchies this one lock partakes in, there is no need for this annotation.

Was this patch driven by a lockdep report?

```
> Signed-off-by: Dhaval Giani <dhaval@linux.vnet.ibm.com>  
> Signed-off-by: Gautham Shenoy R <ego@in.ibm.com>  
>  
>  
> Index: linux-2.6.23-rc1/mm/memcontrol.c  
> ======  
> --- linux-2.6.23-rc1.orig/mm/memcontrol.c 2007-07-30 17:27:24.000000000 +0530  
> +++ linux-2.6.23-rc1/mm/memcontrol.c 2007-07-30 18:43:40.000000000 +0530  
> @@ -501,6 +501,9 @@
```

```
>
> static struct mem_container init_mem_container;
>
> /* lockdep should know about lru_lock */
> +static struct lock_class_key lru_lock_key;
> +
> static struct container_subsys_state *
> mem_container_create(struct container_subsys *ss, struct container *cont)
> {
> @@ -519,6 +522,7 @@
> INIT_LIST_HEAD(&mem->active_list);
> INIT_LIST_HEAD(&mem->inactive_list);
> spin_lock_init(&mem->lru_lock);
> +lockdep_set_class(&mem->lru_lock, &lru_lock_key);
> mem->control_type = MEM_CONTAINER_TYPE_ALL;
> return &mem->css;
> }
```

Containers mailing list

Containers@lists.linux-foundation.org

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