

---

Subject: Re: [RFC][PATCH] Make access to taks's nsproxy liter  
Posted by [Pavel Emelianov](#) on Thu, 09 Aug 2007 07:09:55 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Oleg Nesterov wrote:

> On 08/08, Eric W. Biederman wrote:

>> Oleg Nesterov <[oleg@tv-sign.ru](mailto:oleg@tv-sign.ru)> writes:

>>

>>> On 08/08, Pavel Emelyanov wrote:

>>>> +void switch\_task\_namespaces(struct task\_struct \*p, struct nsproxy \*new)

>>>> +{

>>>> + struct nsproxy \*ns;

>>>> +

>>>> + might\_sleep();

>>>> +

>>>> + ns = p->nsproxy;

>>>> + if (ns == new)

>>>> + return;

>>>> +

>>>> + if (new)

>>>> + get\_nsproxy(new);

>>>> + rcu\_assign\_pointer(p->nsproxy, new);

>>>> +

>>>> + if (ns && atomic\_dec\_and\_test(&ns->count)) {

>>>> + /\*

>>>> + \* wait for others to get what they want from this

>>>> + \* nsproxy. cannot release this nsproxy via the

>>>> + \* call\_rcu() since put\_mnt\_ns will want to sleep

>>>> + \*/

>>>> + synchronize\_rcu();

>>>> + free\_nsproxy(ns);

>>>> + }

>>>> + }

>>> (I may be wrong, Paul cc'ed)

>>>

>>> This is correct with the current implementation of RCU, but strictly speaking,

>>> we can't use synchronize\_rcu() here, because write\_lock\_irq() doesn't imply

>>> rcu\_read\_lock() in theory.

```
void __lockfunc _write_lock(rwlock_t *lock)
```

```
{
```

```
    preempt_disable();
```

```
    rwlock_acquire(&lock->dep_map, 0, 0, _RET_IP_);
```

```
    LOCK_CONTENDED(lock, _raw_write_trylock, _raw_write_lock);
```

```
}
```

preempt\_disable == rcu\_read\_lock() due to

```
#define rcu_read_lock() \
```

```

do { \
    preempt_disable(); \
    __acquire(RCU); \
} while(0)

```

so currently this is enough to write\_lock()

>> But we should be able to do:

```

>>
>> write_lock_irq();
>> rcu_read_lock();
>> muck with other tasks nsproxy.
>> rcu_read_unlock();
>> write_unlock_irq();
>>

```

>> Which would make rcu fine.

>  
> Yes sure. I just meant that the patch looks incomplete. But we didn't  
> hear Paul yet, perhaps I'm just wrong.

>  
>> The real locking we have is that only a task is allowed to modify it's  
>> own nsproxy pointer. Other processes are not.

>>  
>> The practical question is how do we enable other processes to read  
>> a particular tasks nsproxy or something pointed to by it?

>  
> task\_lock(). The only problem we can't take it in do\_notify\_parent(),  
> but if we add read\_lock(tasklist) to sys\_unshare, we can safely access  
> ->parent->nsproxy.

we can safely access parent's nsproxy with this patch like this:

```

rcu_read_lock();
nsproxy = task_nsproxy(p->parent);
BUG_ON(nsproxy == NULL); /* parent should reparent us before exiting nsproxy */
pid_ns = nsproxy->pid_ns;
...
rcu_read_unlock();

```

>  
> Oleg.  
>  
>

---

Containers mailing list  
Containers@lists.linux-foundation.org

