Subject: Re: [PATCH 2.6.24-rc8-mm1 12/15] (RFC) IPC/semaphores: make use of RCU to free the sem_undo_list
Posted by Pierre Peiffer on Thu, 31 Jan 2008 09:52:08 GMT
View Forum Message <> Reply to Message

Serge E. Hallyn wrote:
> Quoting pierre.peiffer@bull.net (pierre.peiffer@bull.net):
>> From: Pierre Peiffer <pierre.peiffer@bull.net>
>>
>> Today, the sem_undo_list is freed when the last task using it exits.
>> There is no mechanism in place, that allows a safe concurrent access to
>> the sem_undo_list of a target task and protects efficiently against a
>> task-exit.
>>
>> That is okay for now as we don't need this.
>>
>> As I would like to provide a /proc interface to access this data, I need
>> such a safe access, without blocking the target task if possible.
>>
>> This patch proposes to introduce the use of RCU to delay the real free of
>> these sem_undo_list structures. They can then be accessed in a safe manner
>> by any tasks inside read critical section, this way:
>>
>> struct sem_undo_list *undo_list;
>> int ret;
>> ...
>> rcu_read_lock();
>> undo_list = rcu_dereference(task->sysvsem.undo_list);
>> if (undo_list)
>>   ret = atomic_inc_not_zero(&undo_list->refcnt);
>> rcu_read_unlock();
>> ...
>> if (undo_list && ret) {
>>   /* section where undo_list can be used quietly */
>>   ...
>> }
>> ...
>
> And of course then
>
> if (atomic_dec_and_test(&undo_list->refcnt))
>   free_semundo_list(undo_list);
>
> by that task.
>

I will precise this too.

>> Signed-off-by: Pierre Peiffer <pierre.peiffer@bull.net>
>
> Looks correct in terms of locking/refcounting.
>
> Signed-off-by: Serge Hallyn <serue@us.ibm.com>
>

Thanks !

--
Pierre

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