Subject: Re: [patch 06/20] [Network namespace] Move the nsproxy NULL affection
Posted by Mishin Dmitry on Mon, 11 Dec 2006 15:22:58 GMT
View Forum Message <> Reply to Message

On Monday 11 December 2006 00:58, dlezcano@fr.ibm.com wrote:
> Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>
>
> ---
>
> kernel/nsproxy.c | 2 +-
> 1 files changed, 1 insertion(+), 1 deletion(-)
>
> Index: 2.6.19-rc6-mm2/kernel/nsproxy.c
> ========== =======
> --- 2.6.19-rc6-mm2.orig/kernel/nsproxy.c
> +++ 2.6.19-rc6-mm2/kernel/nsproxy.c
> @@ -54,10 +54,10 @@ void exit_task_namespaces(struct task_st
> {
> struct nsproxy *ns = p->nsproxy;
> if (ns) {
> + put_nsproxy(ns);
> task_lock(p);
> p->nsproxy = NULL;
> task_unlock(p);
> - put_nsproxy(ns);
> }
> }
This will follow in a race.

--
Thanks,
Dmitry.

Subject: Re: [patch 06/20] [Network namespace] Move the nsproxy NULL affection
Posted by Daniel Lezcano on Mon, 11 Dec 2006 15:46:06 GMT
View Forum Message <> Reply to Message

Dmitry Mishin wrote:
> On Monday 11 December 2006 00:58, dlezcano@fr.ibm.com wrote:
>> Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>
>>
>> ---
>>
>> kernel/nsproxy.c | 2 +-
>> 1 files changed, 1 insertion(+), 1 deletion(-)
>>

>> Index: 2.6.19-rc6-mm2/kernel/nsproxy.c
>> ===
>> --- 2.6.19-rc6-mm2.orig/kernel/nsproxy.c
>> +++ 2.6.19-rc6-mm2/kernel/nsproxy.c
>> @@ -54,10 +54,10 @@ void exit_task_namespaces(struct task_st
>> {
>>   struct nsproxy *ns = p->nsproxy;
>>   if (ns) {
>> + put_nsproxy(ns);
>>   task_lock(p);
>>   p->nsproxy = NULL;
>>   task_unlock(p);
>> - put_nsproxy(ns);
>>   }
>> }
> This will follow in a race.

Yep.

I did that to have a quick fix for the net_ns cleanup.
The problem raised is the p->nsproxy = NULL followed by put_ns_proxy and
put_net_ns. This one will call free_net_ns followed by ip_fib_cleanup
and that will use current_net_ns (which is current->nsproxy->net_ns)
with nsproxy = NULL

I think, the right fix is to have ip_fib_cleanup called with the net_ns
pointer as parameter and in this case current_net_ns needs to be checked
in all the code to ensure it will not be used by the cleanup.

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [patch 06/20] [Network namespace] Move the nsproxy NULL affection
Posted by Mishin Dmitry on Mon, 11 Dec 2006 15:51:39 GMT
View Forum Message <> Reply to Message

On Monday 11 December 2006 18:46, Daniel Lezcano wrote:
> Dmitry Mishin wrote:
> > On Monday 11 December 2006 00:58, dlezcano@fr.ibm.com wrote:
> >> Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>
> >>
> >> ---
> >>
> >> kernel/nsproxy.c | 2 +-

> >> 1 files changed, 1 insertion(+), 1 deletion(-)
> >>
> >> Index: 2.6.19-rc6-mm2/kernel/nsproxy.c
> >> ===
> >> --- 2.6.19-rc6-mm2.orig/kernel/nsproxy.c
> >> +++ 2.6.19-rc6-mm2/kernel/nsproxy.c
> >> @@ -54,10 +54,10 @@ void exit_task_namespaces(struct task_st
> >> {
> >>   struct nsproxy *ns = p->nsproxy;
> >>   if (ns) {
> >> + put_nsproxy(ns);
> >>   task_lock(p);
> >>   p->nsproxy = NULL;
> >>   task_unlock(p);
> >> - put_nsproxy(ns);
> >>   }
> >> }
> > This will follow in a race.
>
> Yep.
>
> I did that to have a quick fix for the net_ns cleanup.
> The problem raised is the p->nsproxy = NULL followed by put_ns_proxy and
> put_net_ns. This one will call free_net_ns followed by ip_fib_cleanup
> and that will use current_net_ns (which is current->nsproxy->net_ns)
> with nsproxy = NULL
>
> I think, the right fix is to have ip_fib_cleanup called with the net_ns
> pointer as parameter and in this case current_net_ns needs to be checked
> in all the code to ensure it will not be used by the cleanup.
Yes, it is. Shame on me, I've missed it when porting. Will be fixed in next version.

--
Thanks,
Dmitry.

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers