
Subject: Re: PATCH -mm] fix create_new_namespaces() return value
Posted by [Cedric Le Goater](#) on Tue, 12 Jun 2007 07:49:10 GMT

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

Badari Pulavarty wrote:

```
>  
>  
> Cedric Le Goater wrote:  
>  
>> The following patch modifies create_new_namespaces() to also use the  
>> errors returned by the copy_*_ns routines and not to systematically  
>> return ENOMEM.  
>>  
>  
> In my initial version, I did same. It doesn't work :(  
>  
> copy_*_ns() routines doesn't return any errors. All they return is NULL  
> in case of a  
> failure + with the exception of copy_mnt_ns, there are no other failure  
> cases.
```

mnt ant uts namespaces return NULL right but this is not true for ipc namespace which returns -ENOMEM. user namespace returns -EINVAL if CONFIG_USER_NS is not selected. I expect pid namespace to do same and net namespace to do even more.

```
> So, there is no way to find out why the copy_*_ns() routines failed from  
> create_new_namespaces().  
> If you really really want to do this, change all copy_*_ns() routines to  
> returns meaningful errors instead of NULL.
```

I will certainly do and change mnt ant uts namespaces to return -ENOMEM but we will still require the patch below.

```
>> Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>  
>> Cc: Serge E. Hallyn <serue@us.ibm.com>  
>> Cc: Badari Pulavarty <pbadari@us.ibm.com>  
>> Cc: Pavel Emelianov <xemul@openvz.org>  
>> Cc: Herbert Poetzl <herbert@13thfloor.at>  
>> Cc: Eric W. Biederman <ebiederm@xmission.com>  
>> ---  
>> kernel/nsproxy.c | 23 ++++++-----  
>> 1 file changed, 17 insertions(+), 6 deletions(-)  
>>  
>> Index: 2.6.22-rc4-mm2/kernel/nsproxy.c  
>> ======  
>> --- 2.6.22-rc4-mm2.orig/kernel/nsproxy.c  
>> +++ 2.6.22-rc4-mm2/kernel/nsproxy.c  
>> @@ -58,30 +58,41 @@ static struct nsproxy *create_new_namesp
```

```

>>     struct fs_struct *new_fs)
>> {
>>     struct nsproxy *new_nsp;
>> +   int err;
>>
>>     new_nsp = clone_nsproxy(tsk->nsproxy);
>>     if (!new_nsp)
>>         return ERR_PTR(-ENOMEM);
>>
>>     new_nsp->mnt_ns = copy_mnt_ns(flags, tsk->nsproxy->mnt_ns, new_fs);
>> -   if (IS_ERR(new_nsp->mnt_ns))
>> +   if (IS_ERR(new_nsp->mnt_ns)) {
>> +       err = PTR_ERR(new_nsp->mnt_ns);
>>     goto out_ns;
>> +   }
>>
>>     new_nsp->uts_ns = copy_utsname(flags, tsk->nsproxy->uts_ns);
>> -   if (IS_ERR(new_nsp->uts_ns))
>> +   if (IS_ERR(new_nsp->uts_ns)) {
>> +       err = PTR_ERR(new_nsp->uts_ns);
>>     goto out_uts;
>> +   }
>>
>>     new_nsp->ipc_ns = copy_ipcs(flags, tsk->nsproxy->ipc_ns);
>> -   if (IS_ERR(new_nsp->ipc_ns))
>> +   if (IS_ERR(new_nsp->ipc_ns)) {
>> +       err = PTR_ERR(new_nsp->ipc_ns);
>>     goto out_ipc;
>> +   }
>>
>>     new_nsp->pid_ns = copy_pid_ns(flags, tsk->nsproxy->pid_ns);
>> -   if (IS_ERR(new_nsp->pid_ns))
>> +   if (IS_ERR(new_nsp->pid_ns)) {
>> +       err = PTR_ERR(new_nsp->pid_ns);
>>     goto out_pid;
>> +   }
>>
>>     new_nsp->user_ns = copy_user_ns(flags, tsk->nsproxy->user_ns);
>> -   if (IS_ERR(new_nsp->user_ns))
>> +   if (IS_ERR(new_nsp->user_ns)) {
>> +       err = PTR_ERR(new_nsp->user_ns);
>>
> Hmm.. copy_user_ns() ? I don't see this in rc4-mm2.

```

it's in the -mm stack.

C.

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