
Subject: Re: [PATCH 2/2] containers: implement namespace tracking subsystem
(v3)

Posted by [serue](#) on Tue, 26 Jun 2007 22:51:04 GMT

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

Quoting Andrew Morton (akpm@linux-foundation.org):

```
> On Wed, 13 Jun 2007 18:01:25 -0500
> "Serge E. Hallyn" <serue@us.ibm.com> wrote:
>
>> @@ -177,6 +184,15 @@ int unshare_nsproxy_namespaces(unsigned long unshare_flags,
>>     if (IS_ERR(*new_nsp)) {
>>         err = PTR_ERR(*new_nsp);
>>         put_nsproxy(old_ns);
>>     goto out;
>> }
>> +
>> + err = ns_container_clone(current);
>> + if (err) {
>>     put_nsproxy(*new_nsp);
>>     put_nsproxy(old_ns);
>> }
>> +
>> +out:
>>     return err;
>> }
>
> I had to fix a reject here: the put_nsproxy(old_ns) has disappeared from
> this code.
```

Thanks Andrew, end result looks good.

-serge

```
> end result:
>
> int unshare_nsproxy_namespaces(unsigned long unshare_flags,
>     struct nsproxy **new_nsp, struct fs_struct *new_fs)
> {
>     int err = 0;
>
>     if (!(unshare_flags & (CLONE_NEWNS | CLONE_NEWUTS | CLONE_NEWIPC |
>             CLONE_NEWUSER)))
>         return 0;
>
>     if (!capable(CAP_SYS_ADMIN))
>         return -EPERM;
>
>     *new_nsp = create_new_namespaces(unshare_flags, current,
```

```
> new_fs ? new_fs : current->fs);
> if (IS_ERR(*new_nsp))
> err = PTR_ERR(*new_nsp);
> goto out;
> }
>
> err = ns_container_clone(current);
> if (err)
> put_nsproxy(*new_nsp);
>
> out:
> return err;
> }
```
