

---

Subject: Re: [PATCH 1/2] SYSCTL: root unregister routine introduced  
Posted by [Stanislav Kinsbursky](#) on Tue, 13 Dec 2011 09:02:52 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> On Mon, 12 Dec 2011 21:50:00 +0300  
> Stanislav Kinsbursky<skinsbursky@parallels.com> wrote:  
>  
>> This routine is required for SUNRPC sysctl's, which are going to be allocated,  
>> processed and destroyed per network namespace context.  
>> IOW, new sysctl root will be registered on network namespace creation and  
>> thus have to unregistered before network namespace destruction.  
>>  
>  
> It's a bit suspicious that such a mature subsystem as sysctl newly  
> needs its internals exported like this. Either a) the net namespaces  
> work is doing something which hasn't been done before or b) it is doing  
> something wrong.  
>  
> So, please explain further so we can confirm that it is a) and not b).  
>

Hello, Andrew.

The goal is to provide an ability to control and modify data by sysctl's in network namespace context. This is done by "net" sysctl's.

But there are two more issues to solve:

- 1) Sysctl's have to be in /proc/sys/sunrpc
- 2) Sysctl's content should be accessible from creator's network context (not current user ones's).

```
>> --- a/kernel/sysctl.c
>> +++ b/kernel/sysctl.c
>> @@ -1701,6 +1701,13 @@ void register_sysctl_root(struct ctl_table_root *root)
>>   spin_unlock(&sysctl_lock);
>> }
>>
>> +void unregister_sysctl_root(struct ctl_table_root *root)
>> +{
>> + spin_lock(&sysctl_lock);
>> + list_del(&root->root_list);
>> + spin_unlock(&sysctl_lock);
>> +}
>> +
>
> This requires the addition of a declaration to include/linux/sysctl.h.
>
> Once that is done and review is complete, I'd suggest that these two
> patches be joined into a single patch, and that patch become part of
```

> whatever patch series it is which needs them.

>

Ok, I'll do so.

--

Best regards,  
Stanislav Kinsbursky