
Subject: Re: Re: [PATCH 1/2] SYSCTL: root unregister routine introduced
Posted by [Glauber Costa](#) on Tue, 13 Dec 2011 10:04:08 GMT

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

On 12/13/2011 02:03 PM, Kinsbursky Stanislav wrote:

>> On 12/13/2011 01:02 PM, Stanislav Kinsbursky wrote:

>>>> 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).

>>>

>> Have you taken a look at how it is done at net/ipv4/sysctl_tcp_ipv4.c ,

>> for instance?

>

> I don't have this file.

> Probably you are talking about net/ipv4/sysctl_net_ipv4.c, don't you?

Yeah, my bad.

>> It manages to handle a per-net sysctl table without touching a single

>> bit at the kernel's core sysctl routines. Not entirely sure if it would

>> fit your use case, but maybe it is worth taking a look.

>>

>> That file achieves both 1) and 2) that you described...

>>

>>
