
Subject: Re: [PATCH 1/2] SYSCTL: root unregister routine introduced
Posted by [ebiederm](#) on Sat, 17 Dec 2011 22:30:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

Stanislav Kinsbursky <skinsbursky@parallels.com> writes:

>> 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 be 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

The sysctl root has nothing to do with what directory the files show up in,
so this should not be an issue.

> 2) Sysctl's content should be accessible from creator's network context (not
> current user's).

Making the sunrpc sysctls per network namespace would seem to address
this. I don't see why you would need a new root to handle this case.

Eric
