Subject: Re: [PATCH net-2.6.25 10/11][INET] Eliminate difference in actions of sysctl and proc handler for co

Posted by Pavel Emelianov on Wed, 05 Dec 2007 09:58:16 GMT

View Forum Message <> Reply to Message

```
David Miller wrote:
> From: Pavel Emelyanov < xemul@openvz.org>
> Date: Tue, 04 Dec 2007 13:16:45 +0300
>> AFAIS the net.ipv4.conf. <dev>, all and default sysctls should
>> work like this when changed (besides changing the value itself):
>>
>> <dev> : optionally do smth else
       : walk devices
>> all
>> default : walk devices
>>
>> The proc handler for net.ipv4.conf.all works like this:
>> <dev> : flush rt cache
>> all : walk devices and flush rt cache
>> default : nothing
>>
>> while the sysctl handler works like this:
>>
>> <dev> : nothing
       : nothing
>> all
>> default: walk devices but don't flush the cache
>> All this looks strange. Am I right that regardless of whatever
>> handler (proc or syscall) is called the behavior should be:
>>
>> <dev> : flush rt cache
       : walk the devices and flush the cache
>> default: walk the devices and flush the cache
>>
>> ?
>>
>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
> Because, basically, nobody (and I really do mean nobody)
> uses the sysctl() method to change these things, what
> people expect is basically going to be the procfs
> access behavior.
OK. Thank you for clarification:)
> And I agree with it.
>
```

> The 'default' influences future settings, it should not modify > existing devices. That's the job of 'all'.

I thought the same, and I saw that this is true for ipv6, but ipv4 works differently: (-- changing default for some sysctls will cause some devices to be changed as well.

I mean - devinet\_copy\_dflt\_conf() copies the changed bit on those devices, that have not this but marked in the "state" field. It is called for such entries as "accept\_redirects", "shared\_media" and many others. But not for "forwarding" one. That's what seemed strange to me. Sorry, that I didn't express the idea more cleanly.

So what's the right behavior -- to propagate the default for all the ctls on all the devices (according to their "state"), not to propagate for all the ctls, or to keep things as they are now?

- > Otherwise why have 'all' and 'default' as two different knobs
- > if they do exactly the same thing? That's pointless.

>

> I've therefore dropped this patch.

>

Thanks, Pavel