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

Posted by [davem](#) on Wed, 05 Dec 2007 10:06:45 GMT

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

From: Pavel Emelyanov <xemul@openvz.org>

Date: Wed, 05 Dec 2007 12:58:16 +0300

> David Miller wrote:

> > 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?

Grrr, good question.

I remember we had all kinds of issues wrt. this which we had to cleanup in IPV6 in particular.

People complained that once a device was loaded and present, you couldn't set the 'default' and expect it to influence the settings.

The user is pretty much screwed in one way or the other. For example:

1) If 'default' propagates to all devices, any specific setting for a device is lost.

2) If 'default' does not propagate, there is no way to have 'default' influence devices which have already been loaded.

I think both behaviors are bad, and the whole problem is that sysctls acting as defaults cannot have sane semantics because devices get loaded before userspace can sanely start making changes to such sysctl

'defaults'.
