

---

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 Thu, 06 Dec 2007 12:31:14 GMT

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

---

Herbert Xu wrote:

> David Miller <davem@davemloft.net> wrote:

>> 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.

>

> Well the way it works on IPv4 currently (for most options) is

> that we'll propagate default settings to a device until either:

>

> 1) the user modifies the setting for that device;

> 2) or that an IPv4 address has been added to the device.

BTW, this is not 100% true. Look, in rtm\_to\_ifaddr()

I see the following code flow:

```
    ipv4_devconf_setall(in_dev);
```

```
    ifa = inet_alloc_ifa();
```

```
    if (ifa == NULL) {
```

```
        /*
```

```
         * A potential indev allocation can be left alive, it stays
```

```
         * assigned to its device and is destroy with it.
```

```
        */
```

```
        err = -ENOBUFS;
```

```
        goto errout;
```

```
    }
```

if we fail to allocate the ifa (hard to happen, but), we will make this device not to accept the default propagation.

If this is a relevant note, I can prepare the patch.

> 2) was done to preserve backwards compatibility as the controls

> were previously only available after address addition and we did

> not propagate default settings in that case..

>

> We could easily extend this so that the default propagation

> worked until the user modified the setting, with an ioctl to  
> revert to the current behaviour for compatibility.  
>  
> Cheers,

---