

---

Subject: Re: [PATCH net-2.6.26 2/6][NETNS][SOCK]: Introduce per-net inuse counters.

Posted by [davem](#) on Fri, 28 Mar 2008 00:38:23 GMT

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

---

From: Eric Dumazet <dada1@cosmosbay.com>

Date: Thu, 27 Mar 2008 21:21:02 +0100

> Given that :

>

> 1) pcounter should really go away from kernel, since Andrew disagree

> with the implementation.

>

> 2) the need to enumerate all protocols in your enum, it seems ... ugly :)

>

> 3) alloc\_percpu(struct net\_prot\_inuse) per net is nice because we dont

> waste memory (if we had to use percpu\_counters for each proto for example)

>

> I suggest to :

>

> 1) not use pcounter anymore

>

> 2) change 'inuse' field to 'inuse\_idx' or 'prot\_num' that is

> automatically allocated at proto\_register time, instead statically at

> compile time.

>

> Just provide a big enough NET\_INUSE\_NR (might depend on IPV6 present or

> not, static or module) to take into account all possible protocols.

>

> struct net\_prot\_inuse {

> int val[NET\_INUSE\_NR];

> };

I agree mostly, and I'll thus hold off on these patches until the details are hashed out.

---