
Subject: [PATCH][PACKET] Remove unneeded packet_socks_nr variable

Posted by [Pavel Emelianov](#) on Wed, 07 Nov 2007 15:32:51 GMT

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

This one is used only under ifdef PACKET_REFCNT_DEBUG in printk and is not needed otherwise. So hide all this stuff under the PACKET_REFCNT_DEBUG.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

```
---
diff --git a/net/packet/af_packet.c b/net/packet/af_packet.c
index 4cb2dfb..e6a96ee 100644
--- a/net/packet/af_packet.c
+++ b/net/packet/af_packet.c
@@ -139,8 +139,30 @@ dev->hard_header == NULL (ll header is added by device, we cannot
control it)
static HLIST_HEAD(packet_sklist);
static DEFINE_RWLOCK(packet_sklist_lock);

+#ifdef PACKET_REFCNT_DEBUG
static atomic_t packet_socks_nr;

+static void packet_sock_inc(void)
+{
+ atomic_inc(&packet_socks_nr);
+}
+
+static void packet_sock_dec(void)
+{
+ atomic_dec(&packet_socks_nr);
+ printk(KERN_DEBUG "PACKET socket %p is free, %d are alive\n",
+ sk, atomic_read(&packet_socks_nr));
+}
+#else
+static inline void packet_sock_inc(void)
+{
+}
+
+static inline void packet_sock_dec(void)
+{
+}
+#endif
+

/* Private packet socket structures. */
```

```
@@ -236,10 +258,7 @@ static void packet_sock_destruct(struct sock *sk)
    return;
}
```

```
- atomic_dec(&packet_socks_nr);
-#ifdef PACKET_REFCNT_DEBUG
- printk(KERN_DEBUG "PACKET socket %p is free, %d are alive\n", sk,
atomic_read(&packet_socks_nr));
-#endif
+ packet_sock_dec();
}
```

```
@@ -1010,7 +1029,7 @@ static int packet_create(struct net *net, struct socket *sock, int
protocol)
    po->num = proto;
```

```
    sk->sk_destruct = packet_sock_destruct;
- atomic_inc(&packet_socks_nr);
+ packet_sock_inc();
```

```
/*
 * Attach a protocol block
```

```
--
```

1.5.3.4

Subject: Re: [PATCH][PACKET] Remove unneeded packet_socks_nr variable

Posted by [Araldo Carvalho de M](#) on Wed, 07 Nov 2007 15:50:04 GMT

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

Em Wed, Nov 07, 2007 at 06:32:51PM +0300, Pavel Emelyanov escreveu:

> This one is used only under ifdef PACKET_REFCNT_DEBUG in

> printk and is not needed otherwise. So hide all this stuff

> under the PACKET_REFCNT_DEBUG.

>

> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Look at sk_refcnt_debug_inc, etc and you'll see a more standard way. I forgot to make this when making all protocol families use sk_prot, even if just partially :-)

- Araldo

Subject: Re: [PATCH][PACKET] Remove unneeded packet_socks_nr variable

Posted by [Araldo Carvalho de M](#) on Wed, 07 Nov 2007 15:53:01 GMT

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

Em Wed, Nov 07, 2007 at 01:50:04PM -0200, Arnaldo Carvalho de Melo escreveu:
> Em Wed, Nov 07, 2007 at 06:32:51PM +0300, Pavel Emelyanov escreveu:
>> This one is used only under ifdef PACKET_REFCNT_DEBUG in
>> printk and is not needed otherwise. So hide all this stuff
>> under the PACKET_REFCNT_DEBUG.
>>
>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
>
> Look at sk_refcnt_debug_inc, etc and you'll see a more standard way. I
> forgot to make this when making all protocol families use sk_prot, even
> if just partially :-)

As a bonus you'll get this information on /proc/net/protocols, removing
'-1' from PACKET column for "sockets".

- Arnaldo

Subject: Re: [PATCH][PACKET] Remove unneeded packet_socks_nr variable
Posted by [Pavel Emelianov](#) on Wed, 07 Nov 2007 16:16:14 GMT

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

Arnaldo Carvalho de Melo wrote:

> Em Wed, Nov 07, 2007 at 01:50:04PM -0200, Arnaldo Carvalho de Melo escreveu:
>> Em Wed, Nov 07, 2007 at 06:32:51PM +0300, Pavel Emelyanov escreveu:
>>> This one is used only under ifdef PACKET_REFCNT_DEBUG in
>>> printk and is not needed otherwise. So hide all this stuff
>>> under the PACKET_REFCNT_DEBUG.
>>>
>>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
>> Look at sk_refcnt_debug_inc, etc and you'll see a more standard way. I
>> forgot to make this when making all protocol families use sk_prot, even
>> if just partially :-)
>
> As a bonus you'll get this information on /proc/net/protocols, removing
> '-1' from PACKET column for "sockets".

Hm... I actually thought about this, but I decided that packet
sockets were not accounted in this way deliberately.

So, shall I break this "compatibility" (-1 in proc) and provide
a packet socket number in this file?

> - Arnaldo

Thanks,

Pavel

Subject: Re: [PATCH][PACKET] Remove unneeded packet_socks_nr variable
Posted by [Arnaldo Carvalho de M](#) on Wed, 07 Nov 2007 16:26:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Em Wed, Nov 07, 2007 at 07:16:14PM +0300, Pavel Emelyanov escreveu:
> Arnaldo Carvalho de Melo wrote:
> > Em Wed, Nov 07, 2007 at 01:50:04PM -0200, Arnaldo Carvalho de Melo escreveu:
> >> Em Wed, Nov 07, 2007 at 06:32:51PM +0300, Pavel Emelyanov escreveu:
> >>> This one is used only under ifdef PACKET_REFCNT_DEBUG in
> >>> printk and is not needed otherwise. So hide all this stuff
> >>> under the PACKET_REFCNT_DEBUG.
> >>>
> >>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
> >> Look at sk_refcnt_debug_inc, etc and you'll see a more standard way. I
> >> forgot to make this when making all protocol families use sk_prot, even
> >> if just partially :-)
> >
> > As a bonus you'll get this information on /proc/net/protocols, removing
> > '-1' from PACKET column for "sockets".
>
> Hm... I actually thought about this, but I decided that packet
> sockets were not accounted in this way deliberately.
>
> So, shall I break this "compatibility" (-1 in proc) and provide
> a packet socket number in this file?

Humm, my bad, the sockets column in /proc/net/protocols doesn't come from prot->socks, it comes from prot->sockets_allocated. But the suggestion for using sk_refcnt_debug_inc stands, it is there for when we want to do what the code in pf_packet does: refcount debugging, in fact that code most probably was copy'n'pasted from other, older protocols.

BTW, IPX also uses this technique, patches are welcome to make it also use common infrastructure :-)

- Arnaldo
