

---

Subject: Re: [PATCH] Compact sk\_stream\_mem\_schedule() code  
Posted by [davem](#) on Tue, 20 Nov 2007 07:22:45 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: Arnaldo Carvalho de Melo <[acme@ghostprotocols.net](mailto:acme@ghostprotocols.net)>  
Date: Mon, 19 Nov 2007 17:30:59 -0200

> Em Mon, Nov 19, 2007 at 03:13:44PM +0300, Pavel Emelyanov escreveu:

> > This function references sk->sk\_prot->xxx for many times.

> > It turned out, that there's so many code in it, that gcc

> > cannot always optimize access to sk->sk\_prot's fields.

> >

> > After saving the sk->sk\_prot on the stack and comparing

> > disassembled code, it turned out that the function became

> > ~10 bytes shorter and made less dereferences (on i386 and

> > x86\_64). Stack consumption didn't grow.

> >

> > Besides, this patch drives most of this function into the

> > 80 columns limit.

> >

> > Signed-off-by: Pavel Emelyanov <[xemul@openvz.org](mailto:xemul@openvz.org)>

>

> I wonder if making it 'const struct proto \*prot = sk->sk\_prot;'

>

> would make any difference.

Such experiments are always useful, but I doubt there will  
be substantial gains in this case.

> Acked-by: Arnaldo Carvalho de Melo <[acme@redhat.com](mailto:acme@redhat.com)>

I've applied the patch, thanks Pavel.

---