
Subject: Re: [PATCH 6/8] Make the sk_clone() lighter
Posted by [davem](#) on Thu, 01 Nov 2007 07:26:06 GMT

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

From: Pavel Emelyanov <xemul@openvz.org>

Date: Wed, 31 Oct 2007 16:54:34 +0300

> The sk_prot_alloc() already performs all the stuff needed by the
> sk_clone(). Besides, the sk_prot_alloc() requires almost twice
> less arguments than the sk_alloc() does, so call the sk_prot_alloc()
> saving the stack a bit.
>
> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
>
> ---
>
> diff --git a/net/core/sock.c b/net/core/sock.c
> index e7537e4..c032f48 100644
> --- a/net/core/sock.c
> +++ b/net/core/sock.c
> @@ -976,8 +976,9 @@ void sk_free(struct sock *sk)
>
> struct sock *sk_clone(const struct sock *sk, const gfp_t priority)
> {
> - struct sock *newsk = sk_alloc(sk->sk_net, sk->sk_family, priority, sk->sk_prot, 0);
> -
> + struct sock *newsk;
> +
> + newsk = sk_prot_alloc(sk->sk_prot, priority, sk->sk_family);
> if (newsk != NULL) {
> struct sk_filter *filter;
>

After we make this change, what will set up newsk->sk_net?

That's part of what sk_alloc() was doing for us, and that's
why we need to pass the extra argument.
