
Subject: Re: [PATCH 5/9] network namespaces: async socket operations

Posted by [Andrey Savochkin](#) on Sat, 23 Sep 2006 13:16:43 GMT

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

On Fri, Sep 22, 2006 at 05:33:56PM +0200, Daniel Lezcano wrote:

> Andrey Savochkin wrote:

> > Non-trivial part of socket namespaces: asynchronous events

> > should be run in proper context.

> >

> > Signed-off-by: Andrey Savochkin <saw@swsoft.com>

> > ---

> > af_inet.c | 10 ++++++++

> > inet_timewait_sock.c | 8 ++++++++

> > tcp_timer.c | 9 ++++++++

> > 3 files changed, 27 insertions(+)

> >

> > --- ./net/ipv4/af_inet.c.venssock-asy Mon Aug 14 17:04:07 2006

> > +++ ./net/ipv4/af_inet.c Tue Aug 15 13:45:44 2006

> > @@ -366,10 +366,17 @@ out_rcu_unlock:

> > int inet_release(struct socket *sock)

> > {

> > struct sock *sk = sock->sk;

> > + struct net_namespace *ns, *orig_net_ns;

> >

> > if (sk) {

> > long timeout;

> >

> > + /* Need to change context here since protocol ->close

> > + * operation may send packets.

> > + */

> > + ns = get_net_ns(sk->sk_net_ns);

> > + push_net_ns(ns, orig_net_ns);

> > +

>

> Is it not a race condition here ? What happens if you have a packet

> incoming during the namespace context switching ?

All asynchronous operations (RX softirq, timers) should set their context explicitly, and can't rely on the current context being the right one (or a valid pointer at all).

Andrey
