Subject: Re: [PATCH 2/3] Lost locking in fl6_sock_lookup
Posted by davem on Thu, 18 Oct 2007 12:14:26 GMT
View Forum Message <> Reply to Message

From: Pavel Emelyanov <xemul@openvz.org>
Date: Thu, 18 Oct 2007 16:11:58 +0400

> YOSHIFUJI Hideaki wrote:
> > In article <47174950.6060409@openvz.org> (at Thu, 18 Oct 2007 15:53:52 +0400), Pavel Emelyanov <xemul@openvz.org> says:
> >
> >> This routine scans the ipv6_fl_list whose update is
> >> protected with the socket lock and the ip6_sk_fl_lock.
> >
> >>   struct ip6_flowlabel *fl = sfl->fl;
> >>   if (fl->label == label) {
> >> +   read_unlock_bh(&ip6_sk_fl_lock);
> >>   fl->lastuse = jiffies;
> >>   atomic_inc(&fl->users);
> >>   return fl;
> >
> > We should increment fl->users within the critical section, shouldn't we?
>
> Not necessary. The users is more than zero (because it is
> linked in the sock's list) so garbage collector won't catch
> it in any way.

Right, we're grabbing an "extra" reference here and only
someone who gets the socket lock (which we have) can unlink
it and thus potentially drop the count to zero.