Subject: [PATCH] Fix again the fl6_sock_lookup() fixed locking
Posted by Pavel Emelianov on Thu, 18 Oct 2007 12:36:58 GMT
View Forum Message <> Reply to Message

YOSHIFUJI fairly pointed out, that the users increment should
be done under the ip6_sk_fl_lock not to give IPV6_FL_A_PUT a
chance to put this count to zero and release the flowlabel.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
Cc: YOSHIFUJI Hideaki <yoshfuji@linux-ipv6.org>

```
diff --git a/net/ipv6/ip6_flowlabel.c b/net/ipv6/ip6_flowlabel.c
index e55ae1a..b12cc22 100644
--- a/net/ipv6/ip6_flowlabel.c
+++ b/net/ipv6/ip6_flowlabel.c
@@ -210,9 +210,9 @@ struct ip6_flowlabel * fl6_sock_lookup(struct sock *sk, __be32 label)
  for (sfl=np->ipv6_fl_list; sfl; sfl = sfl->next) {
   struct ip6_flowlabel *fl = sfl->fl;
   if (fl->label == label) {
-   read_unlock_bh(&ip6_sk_fl_lock);
    fl->lastuse = jiffies;
    atomic_inc(&fl->users);
+   read_unlock_bh(&ip6_sk_fl_lock);
    return fl;
   }
  }
```