
Subject: [PATCH 1/3] Lost locking when inserting a flowlabel in ipv6_fl_list
Posted by [Pavel Emelianov](#) on Thu, 18 Oct 2007 11:51:41 GMT

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

The new flowlabels should be inserted into the sock list under the ip6_sk_fl_lock. This was lost in one place.

This list is naturally protected with the socket lock, but the fl6_sock_lookup() is called without it, so another protection is required.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

```
diff --git a/net/ipv6/ip6_flowlabel.c b/net/ipv6/ip6_flowlabel.c
index 217d60f..8550df2 100644
--- a/net/ipv6/ip6_flowlabel.c
+++ b/net/ipv6/ip6_flowlabel.c
@@ -409,6 +409,16 @@ static int ipv6_opt_cmp(struct ipv6_txoptions *o1, struct ipv6_txoptions
 *o2)
     return 0;
 }

+static inline void fl_link(struct ipv6_pinfo *np, struct ipv6_fl_socklist *sfl,
+ struct ip6_flowlabel *fl)
+{
+ write_lock_bh(&ip6_sk_fl_lock);
+ sfl->fl = fl;
+ sfl->next = np->ipv6_fl_list;
+ np->ipv6_fl_list = sfl;
+ write_unlock_bh(&ip6_sk_fl_lock);
+}
+
int ipv6_flowlabel_opt(struct sock *sk, char __user *optval, int optlen)
{
    int err;
@@ -513,11 +523,7 @@ int ipv6_flowlabel_opt(struct sock *sk, char __user *optval, int optlen)
    fl1->linger = fl->linger;
    if ((long)(fl->expires - fl1->expires) > 0)
        fl1->expires = fl->expires;
- write_lock_bh(&ip6_sk_fl_lock);
- sfl1->fl = fl1;
- sfl1->next = np->ipv6_fl_list;
- np->ipv6_fl_list = sfl1;
- write_unlock_bh(&ip6_sk_fl_lock);
+ fl_link(np, sfl1, fl1);
    fl_free(fl);
```

```
return 0;
```

```
@@ -545,9 +551,7 @@ release:
```

```
}  
}
```

```
- sfl1->fl = fl;  
- sfl1->next = np->ipv6_fl_list;  
- np->ipv6_fl_list = sfl1;  
+ fl_link(np, sfl1, fl);  
return 0;
```

```
default:
```

```
--
```

```
1.5.3.4
```

Subject: Re: [PATCH 1/3] Lost locking when inserting a flowlabel in ipv6_fl_list
Posted by [davem](#) on Thu, 18 Oct 2007 12:15:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Pavel Emelyanov <xemul@openvz.org>
Date: Thu, 18 Oct 2007 15:51:41 +0400

> The new flowlabels should be inserted into the sock list
> under the ip6_sk_fl_lock. This was lost in one place.
>
> This list is naturally protected with the socket lock, but
> the fl6_sock_lookup() is called without it, so another
> protection is required.
>
> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Applied.
