

---

Subject: [PATCH] [NETFILTER] Consolidate nf\_sockopt and compat\_nf\_sockopt v2  
Posted by Pavel Emelianov on Tue, 06 Nov 2007 08:29:18 GMT

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

---

Both lookup the nf\_sockopt\_ops object to call the get/set callbacks from, but they perform it in a completely similar way.

Introduce the helper for finding the ops.

Ported at the top of today's net-2.6 tree to resolve conflict with the patch from Alexey Dobriyan.

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

---

```
diff --git a/net/netfilter/nf_sockopt.c b/net/netfilter/nf_sockopt.c
index 2dfac32..87bc144 100644
--- a/net/netfilter/nf_sockopt.c
+++ b/net/netfilter/nf_sockopt.c
@@ @ -60,46 +60,57 @@ void nf_unregister_sockopt(struct nf_sockopt_ops *reg)
}
EXPORT_SYMBOL(nf_unregister_sockopt);

/* Call get/setsockopt() */
static int nf_sockopt(struct sock *sk, int pf, int val,
-       char __user *opt, int *len, int get)
+static struct nf_sockopt_ops *nf_sockopt_find(struct sock *sk, int pf,
+       int val, int get)
{
    struct nf_sockopt_ops *ops;
-   int ret;

    if (sk->sk_net != &init_net)
-   return -ENOPROTOOPT;
+   return ERR_PTR(-ENOPROTOOPT);

    if (mutex_lock_interruptible(&nf_sockopt_mutex) != 0)
-   return -EINTR;
+   return ERR_PTR(-EINTR);

    list_for_each_entry(ops, &nf_sockopts, list) {
        if (ops->pf == pf) {
            if (!try_module_get(ops->owner))
                goto out_nosup;
+
            if (get) {
-               if (val >= ops->get_optmin
```

```

-     && val < ops->get_optmax) {
-     mutex_unlock(&nf_sockopt_mutex);
-     ret = ops->get(sk, val, opt, len);
+     if (val >= ops->get_optmin &&
+         val < ops->get_optmax)
        goto out;
-   }
} else {
-   if (val >= ops->set_optmin
-       && val < ops->set_optmax) {
-     mutex_unlock(&nf_sockopt_mutex);
-     ret = ops->set(sk, val, opt, *len);
+     if (val >= ops->set_optmin &&
+         val < ops->set_optmax)
        goto out;
-   }
}
module_put(ops->owner);
}
}

- out_nosup:
+out_nosup:
+ ops = ERR_PTR(-ENOPROTOOPT);
+out:
    mutex_unlock(&nf_sockopt_mutex);
- return -ENOPROTOOPT;
+ return ops;
+}
+
+/* Call get/setsockopt() */
+static int nf_sockopt(struct sock *sk, int pf, int val,
+                      char __user *opt, int *len, int get)
+{
+ struct nf_sockopt_ops *ops;
+ int ret;
+
+ ops = nf_sockopt_find(sk, pf, val, get);
+ if (IS_ERR(ops))
+     return PTR_ERR(ops);
+
+ if (get)
+     ret = ops->get(sk, val, opt, len);
+ else
+     ret = ops->set(sk, val, opt, *len);

- out:
    module_put(ops->owner);
    return ret;

```

```

}

@@ -124,51 +135,22 @@ static int compat_nf_sockopt(struct sock *sk, int pf, int val,
    struct nf_sockopt_ops *ops;
    int ret;

- if (sk->sk_net != &init_net)
- return -ENOPROTOOPT;
-
-
- if (mutex_lock_interruptible(&nf_sockopt_mutex) != 0)
- return -EINTR;
-
- list_for_each_entry(ops, &nf_sockopts, list) {
- if (ops->pf == pf) {
- if (!try_module_get(ops->owner))
- goto out_nosup;
-
- if (get) {
- if (val >= ops->get_optmin
- && val < ops->get_optmax) {
- mutex_unlock(&nf_sockopt_mutex);
- if (ops->compat_get)
- ret = ops->compat_get(sk,
- val, opt, len);
- else
- ret = ops->get(sk,
- val, opt, len);
- goto out;
- }
- } else {
- if (val >= ops->set_optmin
- && val < ops->set_optmax) {
- mutex_unlock(&nf_sockopt_mutex);
- if (ops->compat_set)
- ret = ops->compat_set(sk,
- val, opt, *len);
- else
- ret = ops->set(sk,
- val, opt, *len);
- goto out;
- }
- }
- module_put(ops->owner);
- }
+ ops = nf_sockopt_find(sk, pf, val, get);
+ if (IS_ERR(ops))
+ return PTR_ERR(ops);
+

```

```
+ if (get) {
+   if (ops->compat_get)
+     ret = ops->compat_get(sk, val, opt, len);
+   else
+     ret = ops->get(sk, val, ops, len);
+ } else {
+   if (ops->compat_set)
+     ret = ops->compat_set(sk, val, ops, *len);
+   else
+     ret = ops->set(sk, val, ops, *len);
}
- out_nosup:
- mutex_unlock(&nf_sockopt_mutex);
- return -ENOPROTOOPT;

- out:
module_put(ops->owner);
return ret;
}
```

---