

---

Subject: [PATCH 9/12 net-2.6.26] [NETNS]: icmp(v6)\_sk should not pin a namespace.

Posted by [den](#) on Fri, 29 Feb 2008 13:40:55 GMT

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

---

So, change icmp(v6)\_sk creation/disposal to the scheme used in the netlink for rtnl, i.e. create a socket in the context of the init\_net and assign the namespace without getting a reference later.

Also use sk\_release\_kernel instead of sock\_release to properly destroy such sockets.

Signed-off-by: Denis V. Lunev <[den@openvz.org](mailto:den@openvz.org)>

Acked-by: Daniel Lezcano <[dlezcano@fr.ibm.com](mailto:dlezcano@fr.ibm.com)>

---

```
net/ipv4/icmp.c | 12 ++++++-----  
net/ipv6/icmp.c | 11 ++++++-----  
2 files changed, 9 insertions(+), 14 deletions(-)
```

```
diff --git a/net/ipv4/icmp.c b/net/ipv4/icmp.c  
index 7c62a0d..97d97ad 100644  
--- a/net/ipv4/icmp.c  
+++ b/net/ipv4/icmp.c  
@@ -1146,14 +1146,8 @@ static void __exit icmp_exit(void)  
{  
    int i;  
  
- for_each_possible_cpu(i) {  
-     struct sock *sk;  
-  
-     sk = __icmp_sk[i];  
-     if (sk == NULL)  
-         continue;  
-     sock_release(sk->sk_socket);  
- }  
+ for_each_possible_cpu(i)  
+ sk_release_kernel(__icmp_sk[i]);  
    kfree(__icmp_sk);  
    __icmp_sk = NULL;  
}  
@@ -1176,6 +1170,8 @@ int __init icmp_init(void)  
    goto fail;  
  
    __icmp_sk[i] = sk = sock->sk;  
+ sk_change_net(sk, &init_net);  
+  
    sk->sk_allocation = GFP_ATOMIC;
```

```

/* Enough space for 2 64K ICMP packets, including
diff --git a/net/ipv6/icmp.c b/net/ipv6/icmp.c
index 3368f32..7341d79 100644
--- a/net/ipv6/icmp.c
+++ b/net/ipv6/icmp.c
@@ -802,6 +802,8 @@ int __init icmpv6_init(void)
}

__icmpv6_sk[i] = sk = sock->sk;
+ sk_change_net(sk, &init_net);
+
sk->sk_allocation = GFP_ATOMIC;
/*
 * Split off their lock-class, because sk->sk_dst_lock
@@ -831,11 +833,8 @@ int __init icmpv6_init(void)
return 0;

fail:
- for (j = 0; j < i; j++) {
- if (!cpu_possible(j))
- continue;
- sock_release(__icmpv6_sk[j]->sk_socket);
- }
+ for (j = 0; j < i; j++)
+ sk_release_kernel(__icmpv6_sk[j]);

return err;
}
@@ -845,7 +844,7 @@ void icmpv6_cleanup(void)
int i;

for_each_possible_cpu(i) {
- sock_release(__icmpv6_sk[i]->sk_socket);
+ sk_release_kernel(__icmpv6_sk[i]);
}
inet6_del_protocol(&icmpv6_protocol, IPPROTO_ICMPV6);
}
--
```

### 1.5.3.rc5

---

Containers mailing list  
 Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---