Subject: [PATCH] [IPV6]: dst_entry leak in ip4ip6_err. (resend)
Posted by den on Mon, 18 Feb 2008 08:59:38 GMT
View Forum Message <> Reply to Message

The result of the ip_route_output is not assigned to skb. This means that
- it is leaked
- possible OOPS below dereferrencing skb->dst
- no ICMP message for this case

Signed-off-by: Denis V. Lunev <den@openvz.org>
---
 net/ipv6/ip6_tunnel.c | 1 +
 1 files changed, 1 insertions(+), 0 deletions(-)

diff --git a/net/ipv6/ip6_tunnel.c b/net/ipv6/ip6_tunnel.c
index 9031e52..cd94064 100644
--- a/net/ipv6/ip6_tunnel.c
+++ b/net/ipv6/ip6_tunnel.c
@@ -550,6 +550,7 @@ ip4ip6_err(struct sk_buff *skb, struct inet6_skb_parm *opt,
    ip_rt_put(rt);
    goto out;
   }
+ skb2->dst = (struct dst_entry *)rt;
  } else {
   ip_rt_put(rt);
   if (ip_route_input(skb2, eiph->daddr, eiph->saddr, eiph->tos,
--
1.5.3.rc5

--
To unsubscribe from this list: send the line "unsubscribe netdev" in
the body of a message to majordomo@vger.kernel.org
More majordomo info at http://vger.kernel.org/majordomo-info.html