Subject: [PATCH][ICMP]: Dst entry leak in icmp_send host re-lookup code (v2).
Posted by Pavel Emelianov on Wed, 26 Mar 2008 09:25:40 GMT
View Forum Message <> Reply to Message

Commit 8b7817f3a959ed99d7443afc12f78a7e1fcc2063 ([IPSEC]: Add ICMP host relookup support) introduced some dst leaks on error paths: the rt pointer can be forgotten to be put. Fix it bu going to a proper label.

Found after net namespace's lo refused to unregister :) Many thanks to Den for valuable help during debugging.

Herbert pointed out, that xfrm_lookup() will put the rtable in case of error itself, so the first goto fix is redundant.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
Signed-off-by: Denis V. Lunev <den@openvz.org>

```
diff --git a/net/ipv4/icmp.c b/net/ipv4/icmp.c
index a13c074..a944e80 100644
--- a/net/ipv4/icmp.c
+++ b/net/ipv4/icmp.c
@@ -591,7 +591,7 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32 info)
   }

   if (xfrm_decode_session_reverse(skb_in, &fl, AF_INET))
-   goto out_unlock;
+   goto ende;

   if (inet_addr_type(net, fl.fl4_src) == RTN_LOCAL)
    err = __ip_route_output_key(net, &rt2, &fl);
@@ -601,7 +601,7 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32 info)

    fl2.fl4_dst = fl.fl4_src;
    if (ip_route_output_key(net, &rt2, &fl2))
-    goto out_unlock;
+    goto ende;

    /* Ugh! */
    odst = skb_in->dst;
@@ -614,7 +614,7 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32 info)
   }

   if (err)
-   goto out_unlock;
+   goto ende;
```

```
err = xfrm_lookup((struct dst_entry **)&rt2, &fl, NULL,
    XFRM_LOOKUP_ICMP);
```