

---

Subject: [PATCH][ICMP]: Dst entry leak in icmp\_send host re-lookup code.

Posted by [Pavel Emelianov](#) on Tue, 25 Mar 2008 15:40:00 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.

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 ff9a8e6..db231cb 100644
--- a/net/ipv4/icmp.c
+++ b/net/ipv4/icmp.c
@@ -594,11 +594,11 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32 info)
    rt = NULL;
    break;
    default:
-    goto out_unlock;
+    goto ende;
}

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);
@@ -608,7 +608,7 @@ 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;
@@ -621,7 +621,7 @@ 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);
```

---

Subject: Re: [PATCH][ICMP]: Dst entry leak in icmp\_send host re-lookup code.  
Posted by [Herbert Xu](#) on Wed, 26 Mar 2008 03:32:00 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Tue, Mar 25, 2008 at 06:40:00PM +0300, Pavel Emelyanov wrote:  
> 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.  
>  
> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>  
> Signed-off-by: Denis V. Lunev <den@openvz.org>

Thanks for catching this!

```
> diff --git a/net/ipv4/icmp.c b/net/ipv4/icmp.c  
> index ff9a8e6..db231cb 100644  
> --- a/net/ipv4/icmp.c  
> +++ b/net/ipv4/icmp.c  
> @@ -594,11 +594,11 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32  
info)  
>     rt = NULL;  
>     break;  
>     default:  
> -     goto out_unlock;  
> +     goto ende;  
> }
```

I'm not sure about this bit though because xfrm\_lookup is meant  
to free the route on error.

```
>     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);  
> @@ -608,7 +608,7 @@ 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;
> @@ -621,7 +621,7 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32 info)
>     }
>
>     if (err)
> -     goto out_unlock;
> +     goto ende;
```

These ones look good.

Cheers,

--  
Visit Openswan at <http://www.openswan.org/>  
Email: Herbert Xu ~{PmV>HI~} <[herbert@gondor.apana.org.au](mailto:herbert@gondor.apana.org.au)>  
Home Page: <http://gondor.apana.org.au/~herbert/>  
PGP Key: <http://gondor.apana.org.au/~herbert/pubkey.txt>

---

---

Subject: Re: [PATCH][ICMP]: Dst entry leak in icmp\_send host re-lookup code.  
Posted by [den](#) on Wed, 26 Mar 2008 07:46:22 GMT

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

---

On Wed, 2008-03-26 at 11:32 +0800, Herbert Xu wrote:

```
> On Tue, Mar 25, 2008 at 06:40:00PM +0300, Pavel Emelyanov wrote:
> > 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.
> >
> > Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
> > Signed-off-by: Denis V. Lunev <den@openvz.org>
>
> Thanks for catching this!
>
> > diff --git a/net/ipv4/icmp.c b/net/ipv4/icmp.c
> > index ff9a8e6..db231cb 100644
> > --- a/net/ipv4/icmp.c
> > +++ b/net/ipv4/icmp.c
> > @@ -594,11 +594,11 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32
info)
> >     rt = NULL;
```

```
> >     break;
> > default:
> > -   goto out_unlock;
> > +   goto ende;
> > }
>
> I'm not sure about this bit though because xfrm_lookup is meant
> to free the route on error.
```

This is not a problem, it assigns NULL in this case to a pointer.

---

---

Subject: Re: [PATCH][ICMP]: Dst entry leak in icmp\_send host re-lookup code.

Posted by [Pavel Emelianov](#) on Wed, 26 Mar 2008 09:02:34 GMT

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

---

Herbert Xu wrote:

```
> On Tue, Mar 25, 2008 at 06:40:00PM +0300, Pavel Emelyanov wrote:
>> 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.
>>
>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
>> Signed-off-by: Denis V. Lunev <den@openvz.org>
>
> Thanks for catching this!
>
>> diff --git a/net/ipv4/icmp.c b/net/ipv4/icmp.c
>> index ff9a8e6..db231cb 100644
>> --- a/net/ipv4/icmp.c
>> +++ b/net/ipv4/icmp.c
>> @@ -594,11 +594,11 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32
info)
>>     rt = NULL;
>>     break;
>> default:
>> -   goto out_unlock;
>> +   goto ende;
>> }
>
> I'm not sure about this bit though because xfrm_lookup is meant
> to free the route on error.
```

Well, it indeed does, we've missed that fact. Even though it is set to NULL, as Denis pointed out, this hunk is now needed here.

David, plz, disregard this patch, I'll send a new one in a moment.

```
>> 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);
>> @@ -608,7 +608,7 @@ 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;
>> @@ -621,7 +621,7 @@ void icmp_send(struct sk_buff *skb_in, int type, int code, __be32 info)
>> }
>>
>> if (err)
>> - goto out_unlock;
>> + goto ende;
>
> These ones look good.
```

Thanks :)

> Cheers,

---

Subject: Re: [PATCH][ICMP]: Dst entry leak in icmp\_send host re-lookup code.  
Posted by [davem](#) on Wed, 26 Mar 2008 09:06:57 GMT

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

---

From: Pavel Emelyanov <xemul@openvz.org>

Date: Wed, 26 Mar 2008 12:02:34 +0300

> David, plz, disregard this patch, I'll send a new one in a moment.

Ok, but please fix the net-2.6.26 build regression I just reported  
first.

---