Subject: Re: [PATCH][ICMP]: Dst entry leak in icmp_send host re-lookup code (v2).
Posted by Julian Anastasov on Wed, 02 Apr 2008 23:29:55 GMT

View Forum Message <> Reply to Message

Hello,

On Wed, 2 Apr 2008, Herbert Xu wrote:

> On Wed, Apr 02, 2008 at 12:19:06PM +0300, Julian Anastasov wrote:
> >
> > play with saddr=0. In my test setup with 2 interfaces ip_route_input
> > failed because I don't have route to the original destination
> > which is now provided as saddr to ip_route_input. No ICMP was sent
> > to sender while previous kernels send ICMP.
>
> Yes this was an oversight.
>
> [ICMP]: Ensure that ICMP relookup maintains status quo
>
> The ICMP relookup path is only meant to modify behaviour when
> appropriate IPsec policies are in place and marked as requiring
> relookups. It is certainly not meant to modify behaviour when
> IPsec policies don't exist at all.
>
> However, due to an oversight on the error paths existing behaviour
> may in fact change should one of the relookup steps fail.
>
> This patch corrects this by redirecting all errors on relookup
> failures to the previous code path. That is, if the initial
> xfrm_lookup let the packet pass, we will stand by that decision
> should the relookup fail due to an error.

I tested this fix and now ICMP error is sent correctly.
There is mistake in my previous email, I said there is no route but
I have route to my destination, only that ARP resolution fails (after
SNAT which makes the things more funny) and ICMP host unreachable
error should be sent. But it does not matter much, only that NAT
can confuse these xfrm calls but this is out of my knowledge.

Regards

--
Julian Anastasov <ja@ssi.bg>