

RE: [fw-wiz] IPSEC over load-shared T1s (per packet)

Source: <http://www.derkeiler.com/Mailing-Lists/Firewall-Wizards/2003-09/0081.html>

From: Ben Nagy (*ben_at_iagu.net*)

Date: 09/20/03

To: <firewall-wizards@honor.icsalabs.com>

Date: Sat, 20 Sep 2003 05:51:54 +0200

I think this is pretty much solved now, but just for the sake of the archives:

The problem was pretty much as I guessed (just lucky ;).

The packets were being sent over alternating links in strict round-robin, which meant that the ESP packets sometimes arrived out of sequence. The IPsec implementation was dropping all the ones with seq < currentseq, which was causing retransmits in the tunneled TCP sessions.

One fix is to use "per destination" load balancing – but that is bad because if all the traffic is VPN then only one link will get used (only one destination).

What I suggested offlist is to look at either ppp-multilink, or MUX/DE-MUX – both of those will make the link look like one big layer2 pipe, which will fix the problem and preserve sequencing. PPP Multilink is software, and simple. MUX stuff is more complicated but faster and can be more flexible.

I also got queries offlist about the E1/T1 RJ connectors. Yes, I did, OK? I was curious. Ow.

ben

> -----Original Message-----

> From: R. DuFresne [<mailto:dufresne@sysinfo.com>]

> Sent: Friday, September 19, 2003 5:32 PM

> To: Ben Nagy

> Cc: TSimons@Delphi-Tech.com; firewall-wizards@honor.icsalabs.com

>

>

> I might be reading Ben wrong, but, I get the impression he is talking

> about session concerency? A 'sticky' bit in the load

> balancer end, such

> that a session started and sent via one t1 remains directed

Firewall-Wizards: RE: [fw-wiz] IPSEC over load-shared T1s (per packet)

> *through that*
> *path for the remainder of the session?*
>
> *Thanks,*
>
> *Ron DuFresne*
>
> *On Thu, 18 Sep 2003, Ben Nagy wrote:*
>
>> *ObBOFH: One of the T1 RJ connectors must be dirty, which is*
> *causing packet*
>> *corruption. Give both the telco jacks a good clean (licking*
> *them works well)*
>> *and see if that fixes the problem. [1]*
>>
>> *Seriously, I do have a theory ;)*
>>
>> *Does this routing guarantee to preserve sequencing?*
>>
>> *If it's really as you described (packets send one for one*
> *via alternate*
>> *links) then you have some potential problems brewing, I think.*
>>
>> *TCP will "work things out" when packets arrive out of*
> *sequence, but with*
>> *IPSec it's left up to the implementation. One security*
> *concern with most*
>> *crypto things is replay protection. IPSec addresses this by using a*
>> *mandatory sequence number in the ESP header. The receiveing*
> *IPSec doesn't*
>> *_have_ to take any notice, but most do. If your receiving*
> *IPSec has enabled*
>> *replay protection then if one link is going faster half the*
> *packets are*
>> *going to get dropped (sequence number < current).*
>>
>> *This would make your tunneled protocol (say TCP) do the*
> *retransmission thing,*
>> *so it would work itself out eventually, but the speed would*
> *indeed suffer*
>> *horribly.*
>>
>> *See if you can convince your router to preserve "IP flows"*
> *and use the two*
>> *links in a more sensible manner. That might help.*
>>
>>
>>
>> *Best of luck,*
>>
>> *ben*

RE: [fw-wiz] IPSEC over load-shared T1s (per packet)

Firewall-Wizards: RE: [fw-wiz] IPSEC over load-shared T1s (per packet)

> >
> > *PS: Let us know when you work it out? This is an interesting one.*
> >
> > *[1] The RJ's are live, for non-network-engineer types. Not
> enough to kill
> you, but it hurts. :)*
> >
> > > -----Original Message-----
> > > *From: firewall-wizards-admin@honor.icsalabs.com
> > > [mailto:firewall-wizards-admin@honor.icsalabs.com] On Behalf
[...]
> > > We setup IP LOAD-SHARING PER-PACKET on each of the serial
> > > links on both
> > > sides (NOC and Us) in order to get an aggregate 3.0mbit.
> > > PER-PACKET routing
> > > alternates usage of the T1s, one for one...
> > >
> > > Since then, VPN performance has taken a dive. Sniffing out
> > > traffic, ESP
> > > packets are sent 3-4 times before they can be properly decrypted.*

firewall-wizards mailing list
firewall-wizards@honor.icsalabs.com
<http://honor.icsalabs.com/mailman/listinfo/firewall-wizards>