

Re: Strange PPPoe problem

Source: <http://linux.derkeiler.com/Mailing-Lists/Debian/2006-03/msg02642.html>

- *From:* anoop aryal <aaryal@xxxxxxxxxxxxxxxxxxxx>
 - *Date:* Fri, 24 Mar 2006 10:33:07 -0600
-

On Friday 24 March 2006 06:55 am, Jacob S wrote:

On Thu, 23 Mar 2006 14:35:20 -0600

anoop aryal <aaryal@xxxxxxxxxxxxxxxxxxxx> wrote:

On Thursday 23 March 2006 10:58, Jacob S wrote:

-----BEGIN PGP
SIGNED MESSAGE-----
Hash: SHA1

Howdy list,

I recently changed ISPs,
away from static ips on a dsl
line to a
single dynamic ip on
Veriz*n's new Fi*S (fiber
optic) service.
The new service uses PPPoe
– not a problem, or so I
thought – I
have PPPoe on my firewall.

Now, I have used PPPoe
from this very same firewall
on a
different dsl line before and
it worked great. But for
some
reason when I do PPPoe for
the new fiber line only http
traffic
works properly. When
downloading e-mail,
everything is fine
until it tries to download the

Re: Strange PPPoe problem

mail (I see it login, get the number of messages to download, and then it tries to start downloading). At this point the e-mail just hangs until it finally times out. It does not seem to be port-related, as I have setup the e-mail server with port-forwarding rules to allow me to download mail on non-standard ports and it exhibits the same problem. And if I do PPPoe on the provided D-Link router, instead of on my firewall, everything (including e-mail) works great.
<snip>

google PMTU to read about this in more detail, but it seriously sounds like icmp 3/4 packets are being dropped somewhere. if you setup your firewall to allow icmp packets of type 3/4 thru, you should be all set (well, you'd hope so anyway). a set of rules like so should do the trick:

```
-A INPUT -p icmp --icmp-type fragmentation-needed -j ACCEPT
-A OUTPUT -p icmp --icmp-type fragmentation-needed -j ACCEPT
-A FORWARD -p icmp --icmp-type fragmentation-needed -j ACCEPT
```

then, make sure you have the iptutils-ping package installed (not the netkit-ping) and try:

```
ping your.mail.host -c 1 -M do -s 1472
```

and you should get back an icmp reply saying what the mtu should be. subtract 28 from it and try pinging with that size and it should go thru. eg, if the reply says mtu = 1492, try:

```
ping your.mail.host -c 1 -M do -s 1464
```

and it should go thru just fine. if you get a request timeout, that means that some routers are just dropping your packets without an icmp 3/4 message. keep reducing the size of your packet and see if you can get anything thru. read up on PMTU for possible solutions. there are ways to stop automatic PMTU discovery etc.

Re: Strange PPPoe problem

Ok, things are getting stranger here.

I ran the iptables rules you suggested and here's the ping results:

```
# ping longbow.arrowway.com -c 1 -M do -s 1472
PING longbow.arrowway.com (66.252.129.166) 1472(1500) bytes of data.
From pool-71-244-52-50.dllstx.fios.verizon.net (71.244.52.50)
icmp_seq=1 Frag needed and DF set (mtu = 1492)

---- longbow.arrowway.com ping statistics ----
0 packets transmitted, 0 received, +1 errors

# ping longbow.arrowway.com -c 1 -M do -s 1464
PING longbow.arrowway.com (66.252.129.166) 1464(1492) bytes of data.
1472 bytes from longbow.arrowway.com (66.252.129.166): icmp_seq=1 ttl=49
time=163 ms

---- longbow.arrowway.com ping statistics ----
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 163.150/163.150/163.150/0.000 ms
```

So then I added the line

```
pty "/usr/sbin/pppoe -I eth0 -T 80 -m 1464"
```

no, the mtu is 1492. you use 1464 in the ping because you are specifying the ping payload. the total packet size ends up being $1464+28 = 1492$. (which is what you had to begin with.) all is not lost, we know PMTU works on your end provided you leave the three iptable rules mentioned above. it seems like a firewall on the mail server is causing icmp 3/4s from reaching the mail server. we can't do much about the firewall. ie, the mail server replied with packets with size = 1500 (most likely) with DF set, the other endpoint of your DSL sent back an icmp 3/4 message back to the mail server to send smaller packets, the firewall protecting the mail server dropped it and the mail server never knew to send you packets of 1492 or less.

there is a 'hack' that may fix this. try and put this in your ruleset in the *mangle table (assuming your dsl interface is ppp0):

```
-A POSTROUTING -p tcp --tcp-flags SYN,RST SYN -o ppp0 -j TCPMSS
--clamp-mss-to-pmtu
```

or use this from the command line if you don't already have the *mangle section:

```
iptables -t mangle -A POSTROUTING -p tcp --tcp-flags SYN,RST SYN -o ppp0 -j
TCPMSS --clamp-mss-to-pmtu
```

what this does is sets the mss of the initial SYN packet to $pmtu - 40$. the receiver (the mail server in this instance), then, uses the mss to calculate

Re: Strange PPPoe problem

the reply packet size instead of the default of 1500. by doing this it would have created packets of the right size (1492) and therefore doesn't need PMTU to work.

again, the above rule is in addition to the previous three.

hope that works for you.

to /etc/ppp/peers/dsl-provider, but the problem continued. After commenting that line back out (so that no pty... -m declaration had been made in the dsl-provider config), I was able to successfully download one single e-mail from a server. There was only one e-mail in that account and it downloaded like normal. So I sent an e-mail to that account, being that it was on a different server from my normal tests, but that one would not download successfully. So it would seem like it had something to do with the size and speed of the one that downloaded properly.

In short, it's still a no go and I have no clue why. The D-Link router still works great, but pppoe from the firewall doesn't.

Any more clues or suggestions, anyone?

TIA,
Jacob

--

anoop
aaryal@xxxxxxxxxxxxxxxxxxxx

--

To UNSUBSCRIBE, email to debian-user-REQUEST@xxxxxxxxxxxxxxxxxxxx with a subject of "unsubscribe". Trouble? Contact listmaster@xxxxxxxxxxxxxxxxxxxx