router - Podrška #15409

MTU - Maximum Transmission Unit

22.09.2008 16:48 - Ernad Husremović

Status:	Zatvoreno	**Početak:**	22.09.2008
Prioritet:	Nizak	**Završetak:**	
Odgovorna osoba:	Jasmin Beganović	**% završeno:**	100%
Kategorija:		**Procjena vremena:**	0.00 sat
Ciljna verzija:			

Opis

http://forum.openwrt.org/viewtopic.php?id=5554

The MTU is the "Maximum Transmission Unit", or in other words how much you can cram into a single IP frame. Any attempt to send larger frames will incur fragmentation or packet loss. The commonly accepted MTU value is 1500 bytes, defined for ethernet in IEEE802.3 (ethernet spec), however when you're talking about encapsulation such as ppp or vpn, you have to consider the overhead caused by the encapsulation.

As an example, if you PPP over ethernet (PPPoE) the MTU value for the ppp interface will likely be 1492 so that after the PPP the result still fits within the 1500 MTU of the ethernet interface.

If you then decide to run a VPN connection over the PPPoE connection, the MTU value for the VPN interface will be even lower it fits within PPPoE's 1492 MTU.

MTU discovery is done by sending out as large of frames as possible with the "Don't Fragment" set; if an ICMP error is returned, it tries again with a lower MTU value. Due to a combination of fear, paranoia and cluelessness, some people will block ICMP error messages, believing that they are the result of some evil hacker trying to artificially limit their MTU; the ICMP error is never recieved and no changes in frame size are made and the resulting performance is poor and unreliable.

Povezani tiketi:

korelira sa router - Podrška #15705: adsl freezone	**Zatvoreno**	**30.10.2008**

Historija

#1 - 22.09.2008 16:49 - Ernad Husremović

na rama-glas router-u:

```
2479 root       2232 S   /usr/sbin/pppd plugin rp-pppoe.so mtu 1492 mru 1492
```

#2 - 22.09.2008 16:49 - Ernad Husremović

MTU discovery is done by sending out as large of frames as possible with the "Don't Fragment" set; if an ICMP error is returned, it tries again with a lower MTU value

#3 - 22.09.2008 17:01 - Ernad Husremović

Configuring OpenWrt / Network
For all protocol types, you can also specify the MTU by using the mtu option.

#4 - 22.09.2008 17:04 - Ernad Husremović

config network mtu

root@router-wan-sa-1:~# vi /etc/config/network

config 'interface' 'wan'

```
        option 'ifname' 'eth0.1'
        option 'proto' 'pppoe'
        option 'username' 'hsamrae'
        option 'password' 'xxxxxxxxxxxxxx'
        option 'defaultroute' '1'
        option 'ppp_redial' 'persist'
        option 'mtu' '1472'
```

config 'interface' 'freezone'

```
        option 'ifname' 'eth0.1'
        option 'proto' 'pppoe'
        option 'username' 'hsamrae@bihnet'
        option 'password' 'xxxxxxxxxxxxx'
        option 'defaultroute' '0'
        option 'ppp_redial' 'persist'
        option 'mtu' '1472'
```

#5 - 22.09.2008 17:06 - Ernad Husremović

root@router-wan-sa-1:~# ps ax | grep mtu

```
  674 root       2228 S    /usr/sbin/pppd plugin rp-pppoe.so mtu 1472 mru 1472 n
 1117 root       2228 S    /usr/sbin/pppd plugin rp-pppoe.so mtu 1472 mru 1472 n
```

#6 - 22.09.2008 17:38 - Ernad Husremović

test zimbra.rama-glas.com MTU 1472

podesio mtu 1472 kao na sigma-com.net

pokušao 2 x slati poruku od 100 KB (hsamrae@bih.net.ba) na sigma_test@rama-glas.com, i nije prošla, drugi put sam ulovio

```
root@zimbra:~# tail /var/log/mail.log --lines=100000 | grep -i timeout
Sep 22 17:32:06 zimbra postfix/smtpd[22422]: timeout after DATA from mta1.bih.net.ba[195.222.33.153]
```

čim sam treći put ponovo oborio pristup direktno zimbri
poruka je stigla na zimbru

#7 - 22.09.2008 17:44 - Ernad Husremović

čitajući ticket 15316 / 28 moja je pretpostavka da bihnetov server pokušava neuspješno uraditi MTU discovery, ali neki uređaj pravi probleme

#8 - 22.09.2008 17:49 - Ernad Husremović

http://www.mynetwatchman.com/kb/ADSL/pppoemtu.htm

ovdje se ominje 1454 MTU kao optimalan

By contrast, the protocol overhead using a 1454 byte MTU is 16.20%.

Although not a staggering difference, using a lower MTU actually reduces ATM overhead by about 0.6% and will thus yield a corresponding increase in user thoroughput: .06% * 1.5Mbps = ~+90Kbps

If you want to understand the details of exactly why overhead is lower, read on:

http://www.mynetwatchman.com/kb/ADSL/pppoemtu.htm
...

#9 - 22.09.2008 18:07 - Ernad Husremović

- Odgovorna osoba promijenjeno iz Ernad Husremović u Jasmin Beganović

- Prioritet promijenjeno iz Normalan u Urgentno

MTU 1454

pokušao postaviti ovaj parametar na rama-glas, ali nakon restarta router mi se više nije javio

isto ovo pokušao na router-wan-sa-1 i sve radi bez problema

#10 - 22.09.2008 18:12 - Ernad Husremović

veoma veoma interesantne rezultate sam dobio nakon što sam podesio ovaj parametar, pustio sam kopiranje sa internet host-a na officesa:

root@bring:~# scp vps.bring.out.ba_usr_var.tar.gz root@officesa.sigma-com.net:/root

dobio sam brzinu koja je oscilirala od 90 KB/sec do 190 KB/sec što je odlično u odnosu na rezultate koje sam ranije imao

#11 - 22.09.2008 18:25 - Ernad Husremović

root@officesa.sigma-com.net's password:

```
vps.bring.out.ba_usr_var.tar.gz               100%  121MB 152.2KB/s   13:32
```

#12 - 22.09.2008 22:29 - Ernad Husremović

pogriješio sam, scp officesa -> vps.bring.out.ba je relevantan za upload

root@router-back:~# scp root vps.bring.out.ba_usr_var.tar.gz root@vps:/root

```
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'vps,209.40.203.155' (RSA) to the list of known hosts.
root@vps's password:
root: No such file or directory
vps.bring.out.ba_usr_var.tar.gz               100%  121MB  26.7KB/s 1:17:16
```

#13 - 22.09.2008 22:39 - Ernad Husremović

http://www.bhtelecom.ba/uploads/media/Uputstvo_za_ADSL.pdf

Parametar „Maximum Transfer Unit (MTU)“ se mora promijeniti sa početne
vrijednosti **1492** na vrijednost **1460**.

#14 - 22.09.2008 23:15 - Ernad Husremović

na našim fwbuilder postavkama router-wan-rg-2 i router-wan-sa-1 na firewall setting je checkirano **Clamp MSS to MTU**

http://lartc.org/howto/lartc.cookbook.mtu-mss.html

#15 - 22.09.2008 23:17 - Ernad Husremović

As explained above, Path MTU Discovery doesn't work as well as it should anymore. If you know for a fact that a hop somewhere in your network has a limited (<1500) MTU, you cannot rely on PMTU Discovery finding this out.

Besides MTU, there is yet another way to set the maximum packet size, the so called Maximum Segment Size. This is a field in the TCP Options part of a SYN packet.

Recent Linux kernels, and a few PPPoE drivers (notably, the excellent Roaring Penguin one), feature the possibility to 'clamp the MSS'.

The good thing about this is that by setting the MSS value, you are telling the remote side unequivocally 'do not ever try to send me packets bigger than this value'. No ICMP traffic is needed to get this to work.

The bad thing is that it's an obvious hack - it breaks 'end to end' by modifying packets. Having said that, we use this trick in many places and it works like a charm.

In order for this to work you need at least iptables-1.2.1a and Linux 2.4.3 or higher. The basic command line is:

```
# iptables -A FORWARD -p tcp --tcp-flags SYN,RST SYN -j TCPMSS  --clamp-mss-to-pmtu
```

This calculates the proper MSS for your link. If you are feeling brave, or think that you know best, you can also do something like this:

```
# iptables -A FORWARD -p tcp --tcp-flags SYN,RST SYN -j TCPMSS --set-mss 128
```

This sets the MSS of passing SYN packets to 128. Use this if you have VoIP with tiny packets, and huge http packets which are causing chopping in your voice calls.

#16 - 22.09.2008 23:52 - Ernad Husremović

MTU issues with tunnel: frag needed, --set-mss, --clamp-mss-to-ptmu

#17 - 23.09.2008 07:54 - Jasmin Beganović

re: note-10

ja sam se jutros normalno konektovao i odradio reboot routera, sve je normalno podigao, vidim da je MTU 1454, napravio sam tunel sa vps-a u flushovao mailq

#18 - 23.09.2008 10:31 - Ernad Husremović

bjasko nisi mi jasan vezano za ramaglas, jeste li morati ugasi/upali router uraditi ?

#19 - 23.09.2008 10:44 - Ernad Husremović

fwbuilder, icmp, mtu parametri

dva su parametra fwbuilder-a interesantna vezano za ovaj problem fragmentacije saobraćaja

1. clamp-mss-to-pmtu
2. icmp saobraćaj koji može biti potrban da pošaljiocu ip paketa poruke fragmentacije (MTU discovery is done by sending out as large of frames as possible with the "Don't Fragment" set; if an ICMP error is returned, it tries again with a lower MTU value)

#20 - 23.09.2008 11:26 - Ernad Husremović

```
(11:06:49) bjasko: meni se je jutros router normalno javio i ja ga resetovao neznam zašto je tebe odbijao
(11:07:04) hernad: hoćeš da kažeš da ga nisi resetovao ?
(11:08:23) bjasko: jesam
(11:08:44) hernad: pa ne kontam šta mi to govoriš
(11:08:54) hernad: ja ti kažem da postoje problemi sa tim router-om kod reboot-a
(11:09:04) hernad: ja sam ga juče jednom resetovao, i sve je bilo ok
(11:09:15) hernad: u roku od 30-sec se ponovo javio i direktno i preko vpn-a
(11:09:21) hernad: a kod drugog reseta se više nije javio
(11:09:25) hernad: otvori za ovo ticket
(11:10:51) bjasko: ok
```

#21 - 23.09.2008 11:26 - Ernad Husremović

ispravka: reset => reboot

#22 - 24.09.2008 02:41 - Ernad Husremović

pokušaću na router-u staviti ovaj mtu

```
config interface        wan
        option ifname   "eth0.1"
        option proto    "pppoe"
...
        option defaultroute "1"
        option ppp_redial "persist"
        option mtu "1000"

config interface        freezone
        option ifname   "eth0.1"
        option proto    "pppoe"
...
        option defaultroute "0"
        option ppp_redial "persist"
        option mtu "1000"
```

#23 - 24.09.2008 02:50 - Ernad Husremović

međutim, to ne pije vode, nisam mogao isporučiti sa vps.bring.out.ba tu poštu

čak sam i ručno naveo

root@router-wan-rg-2:~# ifconfig ppp0 mtu 1000

root@router-wan-rg-2:~# ifconfig ppp0

```
ppp0      Link encap:Point-to-Point Protocol
          inet addr:92.36.147.124  P-t-P:92.36.128.1  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1000  Metric:1
          RX packets:849 errors:0 dropped:0 overruns:0 frame:0
          TX packets:783 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:3
          RX bytes:127950 (124.9 KiB)  TX bytes:240865 (235.2 KiB)
```

i fakat se ovaj parametar izgleda ignoriše

#24 - 24.09.2008 02:51 - Ernad Husremović

mailq na vps bring.out.ba:

```
D9F7E3A70071  2805671 Wed Sep 24 00:37:53  root@vps.bring.out.ba
```

```
(conversation with adsl.rama-glas.com[92.36.147.124] timed out while sending message body)
                                        sigma_test@rama-glas.com
```

#25 - 24.09.2008 02:53 - Ernad Husremović

vratio sam stanje firewall-a tako da ponovo odbija direktan prijem preko adsl.rama-glas.com

#26 - 24.09.2008 02:58 - Ernad Husremović

podesio ssh tunel

root@bring:~# ssh -f root@adsl.rama-glas.com -L 9000:192.168.55.9:25 -N

ispraznio queue

#27 - 09.10.2008 16:44 - Ernad Husremović

kontaktirao me je merim iz bhtelecom-a, i preporučio da podesim mtu i na strani LAN-a na router-u, mislim da je dovoljno da setumem br-lan interfejs:

root@router-wan-rg-2:~# ifconfig br-lan mtu 1000

root@router-wan-rg-2:~# ifconfig

```
br-lan    Link encap:Ethernet  HWaddr 00:1D:7E:55:6D:A6
          inet addr:192.168.55.254  Bcast:192.168.55.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1000  Metric:1
          RX packets:3258157 errors:0 dropped:0 overruns:0 frame:0
          TX packets:3585759 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:557455202 (531.6 MiB)  TX bytes:2699788162 (2.5 GiB)

eth0      Link encap:Ethernet  HWaddr 00:1D:7E:55:6D:A6
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:8979382 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8271486 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:3550533095 (3.3 GiB)  TX bytes:3349610952 (3.1 GiB)
          Interrupt:4

eth0.0    Link encap:Ethernet  HWaddr 00:1D:7E:55:6D:A6
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:3258163 errors:0 dropped:0 overruns:0 frame:0
          TX packets:3585759 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:570488130 (544.0 MiB)  TX bytes:2714131198 (2.5 GiB)

eth0.1    Link encap:Ethernet  HWaddr 00:1D:7E:55:6D:A6
          UP BROADCAST RUNNING MULTICAST  MTU:1000  Metric:1
          RX packets:5721221 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4685126 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:2817707199 (2.6 GiB)  TX bytes:554371966 (528.6 MiB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:1 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:89 (89.0 B)  TX bytes:89 (89.0 B)

ppp0      Link encap:Point-to-Point Protocol
          inet addr:92.36.152.116  P-t-P:92.36.128.1  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1000  Metric:1
          RX packets:341948 errors:0 dropped:0 overruns:0 frame:0
          TX packets:251353 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:3
          RX bytes:165798276 (158.1 MiB)  TX bytes:29659734 (28.2 MiB)

ppp1      Link encap:Point-to-Point Protocol
          inet addr:10.1.248.243  P-t-P:10.1.0.1  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1000  Metric:1
          RX packets:3 errors:0 dropped:0 overruns:0 frame:0
          TX packets:3 errors:0 dropped:0 overruns:0 carrier:0
```

```
          collisions:0 txqueuelen:3
          RX bytes:114 (114.0 B)  TX bytes:54 (54.0 B)

tun0      Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
          inet addr:10.8.0.180  P-t-P:10.8.0.1  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1500  Metric:1
          RX packets:149 errors:0 dropped:0 overruns:0 frame:0
          TX packets:99 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:12316 (12.0 KiB)  TX bytes:14947 (14.5 KiB)
```

#28 - 09.10.2008 17:01 - Ernad Husremović

otvorio sam firewall na router-u rama-glas.com

hernad@nmraka-1:~$ telnet adsl.rama-glas.com 25

```
Trying 92.36.152.116...
Connected to adsl.rama-glas.com.
Escape character is '^]'.
220 zimbra.rama-glas.com ESMTP Postfix
```

#29 - 09.10.2008 17:10 - Ernad Husremović

testirao sa vps.bring.out.ba (američki host) slanje kratke poruke

Oct 9 15:08:54 bring postfix/smtp[14021]: 51D163A70072: to=<sigma_test@rama-glas.com>, relay=adsl.rama-glas.com[92.36.152.116]:25, delay=10, delays=0.53/0.01/0.71/9, dsn=2.0.0, status=sent (250 2.0.0 Ok: queued as 8BD7EF80240)

idemo na dugu poruku:

root@bring:~# mail sigma_test@rama-glas.com -s "duga poruka" < test.txt

#30 - 09.10.2008 17:14 - Ernad Husremović

hej izgleda da je fakat prošla i ova poruka bez problema

root@bring:~# tail /var/log/mail.log --lines=1000 | grep "size=" | grep root@vps.bring.out.ba

```
Oct  9 15:07:53 bring postfix/qmgr[14014]: 471DE3A70072: from=<root@vps.bring.out.ba>, size=336, nrcpt=1 (queu
e active)
Oct  9 15:08:44 bring postfix/qmgr[14014]: 51D163A70072: from=<root@vps.bring.out.ba>, size=2805693, nrcpt=1 (
queue active)
```

#31 - 09.10.2008 17:15 - Ernad Husremović

znači smanjenje mtu-a na router-u na LAN strani pije vode:

```
root@router-wan-rg-2:~# ifconfig br-lan mtu 1000
```

#32 - 09.10.2008 17:25 - Ernad Husremović

idem sada na zimbru rama-glas-ovu:

vidim testove koje radi neven:

root@zimbra:~# tail /var/log/mail.log --lines=1000 | grep "from.*qss"

```
Oct  9 17:18:31 zimbra postfix/smtpd[32261]: connect from qss-server1.qss.ba[195.222.57.201]
Oct  9 17:18:40 zimbra postfix/qmgr[2807]: E2B60F80240: from=<neven@qssbh.com>, size=2003, nrcpt=1 (queue acti
ve)
Oct  9 17:18:40 zimbra postfix/smtpd[32261]: disconnect from qss-server1.qss.ba[195.222.57.201]
Oct  9 17:18:41 zimbra postfix/qmgr[2807]: 1F09DF80243: from=<neven@qssbh.com>, size=2657, nrcpt=1 (queue acti
ve)
```

test-2 sam primio, prvi test nisam. radi se o kratkim porukama

#33 - 09.10.2008 17:26 - Ernad Husremović

međutim, ja sam poslao testnu poruku sa sigma-com.net a mi koristimo bihnet-ov stmp server .. poruka opet nije prošla :(

#34 - 09.10.2008 17:28 - Ernad Husremović

to je ova konekcija

root@zimbra:~# tail /var/log/mail.log --lines=100000 | grep "mta.*bih.net"

```
Oct  9 17:16:22 zimbra postfix/smtpd[31260]: connect from mta2.bih.net.ba[195.222.33.154]
Oct  9 17:16:22 zimbra postfix/smtpd[31260]: B6226F801E8: client=mta2.bih.net.ba[195.222.33.154]
```

ali će to garant nakon 10-tak minuta biti već poznati timeout

#35 - 09.10.2008 17:30 - Ernad Husremović

ispravka note-27: nije merim i nije bhtelecom nego neven iz qss-a

#36 - 09.10.2008 17:32 - Ernad Husremović

nevenova poruka "telnet test 6" je prošao

#37 - 09.10.2008 17:34 - Ernad Husremović

prošle su i poruke:

- test1 (Sun One 6.1), bhtelecom (from gljiva@bih.net.ba)
- test3 (from neven@narkomanija.ba)

#38 - 09.10.2008 17:35 - Ernad Husremović

moja testna poruka sigma-com.net via bihnet je velika cca 350 KB.

#39 - 09.10.2008 17:36 - Ernad Husremović

i evo ga kako sam i očekivao ova poruka nije prošla

root@zimbra:~# tail /var/log/mail.log --lines=100000 | grep "timeout"

```
Oct  9 17:33:02 zimbra postfix/smtpd[31260]: timeout after DATA from mta2.bih.net.ba[195.222.33.154]
```

#40 - 09.10.2008 17:45 - Ernad Husremović

hm moj sljedeći test poruke od 2,8 MB sa vps.bring.out.ba nije prošao

root@bring:~# mailq

```
-Queue ID- --Size-- ----Arrival Time---- -Sender/Recipient-------
A16BB3A70072  2805697 Thu Oct  9 15:37:38  root@vps.bring.out.ba
(conversation with adsl.rama-glas.com[92.36.152.116] timed out while sending message body)
                                          sigma_test@rama-glas.com
```

znači MTU na router-u ipak ne daje željene rezultate

#41 - 09.10.2008 17:46 - Ernad Husremović

u međuvremenu nevenovi testovi sa attachmentima su prošli

- test9 4.3 MB
- test10 320 KB

#42 - 09.10.2008 17:50 - Ernad Husremović

evo ih u u log-u:

root@zimbra:~# tail /var/log/mail.log --lines=10000 | grep gljiva

```
Oct  9 17:31:57 zimbra postfix/qmgr[2807]: D4E4DF80246: from=<gljiva@bih.net.ba>, size=1610, nrcpt=1 (queue ac
tive)
Oct  9 17:31:58 zimbra postfix/qmgr[2807]: 8C3CCF80248: from=<gljiva@bih.net.ba>, size=2251, nrcpt=1 (queue ac
tive)
Oct  9 17:36:13 zimbra postfix/qmgr[2807]: 826A6F80247: from=<gljiva@bih.net.ba>, size=6100723, nrcpt=1 (queue
 active)
Oct  9 17:36:14 zimbra postfix/qmgr[2807]: 2DB7AF80248: from=<gljiva@bih.net.ba>, size=6101300, nrcpt=1 (queue
 active)
```

```
Oct  9 17:36:27 zimbra postfix/qmgr[2807]: 691ABF80247: from=<gljiva@bih.net.ba>, size=519496, nrcpt=1 (queue
active)
Oct  9 17:36:28 zimbra postfix/qmgr[2807]: B08FEF80248: from=<gljiva@bih.net.ba>, size=520202, nrcpt=1 (queue
active)
```

koliko mogu upratiti nisu svi prošli ni od gljiva, barem ne do sada

#43 - 09.10.2008 17:52 - Ernad Husremović

u međuvremenu se nakupilo timeout-a sa raznih servera

root@zimbra:~# tail /var/log/mail.log --lines=10000 | grep timeout

```
Oct  9 17:33:02 zimbra postfix/smtpd[31260]: timeout after DATA from mta2.bih.net.ba[195.222.33.154]
Oct  9 17:36:29 zimbra postfix/smtpd[32261]: timeout after DATA from smtpauth12.prod.mesa1.secureserver.net[64
.202.165.35]
Oct  9 17:43:09 zimbra postfix/smtpd[2799]: timeout after DATA from smtpauth12.prod.mesa1.secureserver.net[64.
202.165.35]
Oct  9 17:51:14 zimbra postfix/smtpd[31260]: timeout after DATA from mta1.bih.net.ba[195.222.33.153]
```

ništa ovo radi ko i prije - tačnije **ne radi** :(

#44 - 09.10.2008 18:19 - Ernad Husremović

to neven@qss.ba

```
Nevene sve stvari vezan za ovaj možete postirati dodavanjem komentara na  http://redmine.bring.out.ba/issues/s
how/15409

bhtelecom/xxxxxxx

ja ću sada vratiti stare postavke pošto pošta kao što vidite zaglavljuje,

pozdrav
Ernad
```

oborio 25 port izvana

hernad@nmraka-1:~$ telnet adsl.rama-glas.com 25

```
Trying 92.36.152.116...
telnet: Unable to connect to remote host: Connection refused
```

#45 - 09.10.2008 18:20 - Ernad Husremović

vps bring out-ba vratio ssh tunel

```
root@bring:~# vi /etc/postfix/transport
root@bring:~# postmap /etc/postfix/transport
root@bring:~# invoke-rc.d postfix restart
 * Stopping Postfix Mail Transport Agent postfix                          [ OK ]
 * Starting Postfix Mail Transport Agent postfix                          [ OK ]
root@bring:~# postfix flush
```

#46 - 09.10.2008 18:22 - Ernad Husremović

evo i druge duge prouke sa vps-2, pošta prošla preko tunela

#47 - 09.10.2008 18:24 - Ernad Husremović

interesantno:

neven je prilikom testa imao problem da pristupi našem redmine-u http://redmine.bring.out.ba

ping ne radi, kada ide preko bihnetovog linka (dns dobro resolvira, ali ping nema replay-a) ??!

preko smartnet linka je normalno pristupio redmine-u ?!?

fakat ne kontam šta je to sa bihnetovom infrastrukturom

#48 - 09.10.2008 18:28 - Ernad Husremović

google kaže:

```
ip_no_pmtu_disc = 0/false means that we DO want MTU discovery.
ip_no_pmtu_disc = 1/true means that we DON't want MTU discovery.
```

#49 - 09.10.2008 18:30 - Ernad Husremović

- Odgovorna osoba promijenjeno iz Jasmin Beganović u bhtelecom bihnet

- % završeno promijenjeno iz 0 u 50

dodijeliću zadatak nevenu, da bi primao email-ove

#50 - 09.10.2008 18:31 - Ernad Husremović

router-rg-2

root@router-wan-sa-1:~# cat /proc/sys/net/ipv4/ip_no_pmtu_disc

```
0
```

#51 - 09.10.2008 18:32 - Ernad Husremović

root@zimbra:~# cat /proc/sys/net/ipv4/ip_no_pmtu_disc

```
0
```

pošto je zimbra openvz sesija, može biti bitno podešenje openvz host-a

root@rg-10:~# cat /proc/sys/net/ipv4/ip_no_pmtu_disc

```
0
```

#52 - 09.10.2008 18:33 - Ernad Husremović

ovo je sve ok mtu discovery se vrši na svim hostovima.

#53 - 09.10.2008 18:38 - Ernad Husremović

neven via email

```
Pozdrav,

zamolio bih vas samo da uradite slijedece prije nego vratite stare postavke:
- podesiti lan ethernet mtu na linksys-u na 1460
- podesiti pppoe mtu na 1460
!! - podesiti mtu na zimbri na 1460

Pa bi probali jos par testova. Ovo (s zimbrom) definitivno nije rjesenje, nego workaround.

Nastavak slijedi...

Neven
```

već sam zatvorio postavke kada sam primio email, ali moram reći da ne očekujem da će ovo uticati na stvar naime ja sam forsiranjem mtu-a 1000 na strani pošiljaoca došao do toga da je komunikacija mimo tunela proradila

svejedno sutra ovo možemo sve probati ako mislite da ipak može biti bitno za rješenje problema.

Međutim, testiranje bih ostavio tek iza 15-16:00 (možemo se naknadno dogovoriti) jer korisnicima ovim testovima praktično onemogućavamo mail.

#54 - 09.10.2008 18:49 - Ernad Husremović

redmine nije nije dopuštao slanje prema vanjskim adresama, podesio sam to, sada bi neven trebao dobijati notifikacijske mailove

#55 - 09.10.2008 18:50 - Ernad Husremović

log kaže da je primio :)

root@mail-gw-10:~# tail /var/log/mail.log | grep neven

```
Oct  9 18:49:35 mail-gw-10 postfix/smtp[14396]: 52EB21A688FF: to=<neven@qss.ba>, relay=out.mail.bih.net.ba[195
.222.33.151]:25, delay=0.63, delays=0.13/0/0.25/0.24, dsn=2.5.0, status=sent (250 2.5.0 Ok.)
```

#56 - 09.10.2008 20:24 - Ernad Husremović

neven via email

```
Pozdrav,
redmine@bring.out.ba je non-reply adresa, na redmine se postira isključivo web interfejsa.
E sad taj timeout ... pitaj me nešto lakše ... da nije mtu :) ?!

da bring.out.ba jeste na google-ovim serverima,  eksperimentišemo sa google-om s obzirom da nas bihnet smtp ze
za non-stop :)

hernad@bring.out.ba je npr validna adresa.

-- "Neven Randic" <neven@qssbh.com> wrote:

> Pozdrav,
>
> problem je MTU na bihnet mrezi. pisem detaljniji izvjestaj.
>
> Note 1: kod postiranja na redmine, sesija time-outira (i s bihnet-a i
> s smartneta), tako da cu dalje izvjestaje slati e-mailom
>
> Note 2: bring.out.ba MX ukazuje na google
>
>  Thu 2008-10-09 18:50:34: [2020:1] *  P=010 D=bring.out.ba TTL=(0)
> MX=[aspmx5.googlemail.com] {74.125.45.27}  Thu 2008-10-09 18:50:34:
> [2020:1] *  P=010 D=bring.out.ba TTL=(0) MX=[aspmx4.googlemail.com]
> {66.249.93.27}  Thu 2008-10-09 18:50:34: [2020:1] *  P=010
> D=bring.out.ba TTL=(0)
>
> account redmine ne postoji pa necu slati mailove tu adresu vec na
> ovu.
>
> Pozdrav,
>
> Neven
```

#57 - 09.10.2008 20:27 - Ernad Husremović

neven via mail, moj odgovor

```
ok, zadatak ću proslijediti kolegi jasminu beganoviću, s obzirom da je on inače zadužen za ove poslove.
Biće Vam na raspolaganju.

----- "Neven Randic" <neven@qssbh.com> wrote:

> jos jednom za svaki slucaj :)
>
> Pozdrav, problem je definitivno mtu na mrezi, pisem detaljniji
> izvjestaj
>
> dogovoreno, testovi sutra u 16h?
>
> ako mozemo uporediti packet capture to ce biti definitivni dokaz
> mrezama da nesto nije u redu. Napisat cu malo opsirniji izvjestaj za
> BIHNET pa cu vam proslijediti.
>
> Neven, neven@qss.ba , 061 890 318
```

#58 - 09.10.2008 20:28 - Ernad Husremović

- Odgovorna osoba promijenjeno iz bhtelecom bihnet u Jasmin Beganović

jale, znači sutra iza 16:00 trebaš biti na raspolaganju nevenu

#59 - 09.10.2008 20:30 - Ernad Husremović

i na kraju ovo je izvještaj koj je neven poslao bihnetu:

```
Pozdrav, QSS je zaduzen za support MTA (Sun JES portal sistema), tako da cu ja prakticno prebaciti problem na mreze
(drugi sektor unutar BIHNET-a).
u nastavku vam saljem izvjestaj koji sam im poslao, koji se smatra confidential, te molim da ga ne proslijedju jete dalje. Na rjesavanju problema cu ostati ukljucen dok se isti ne rijesi uspjesno.

mail za bihnet:

Pozdrav,

s kolegom Ernadom [ernad.husremovic@sigma-com.net] koji je zaduzen da administraciju rama-glas.com servera smo uradili gomilu
testova prilikom slanja i utvrdili da cesto kod velikih e-mail poruka dolazi do time-out-a na SMTP-u. Logovi (smtp+packet capture) su prikupljeni, analizirani.

Paketi prolaze djelimicno do servera (av firewall-a tj. SGS-ova), dosta checksum errora. "MUST FRAGMENT ICMP" paketi ne prolaze! s odredjenih lokacija unutar BIHNET mreze.
Prolaznost malih testnih e-mail poruka 100%. Prolaznost velikih testnih poruka: 25%

Konekcije s drugih mreza prolaze (HT, Smartnet) bez gubitaka ICMP paketa prema ADSL-u sto sugerise da je problem "negdje izmedju",
znaci ne na ADSL/DSLAM-u niti na mail sistemu vec na nekom od rutera.

Problem scope: BIHNET users with non-standard MTU. (MPLS MTU problem mozda?, predlazem tcp mss adjust kao potencijalno rjesenje)

Temporary workaround koji je predlozen korisniku: reduce MTU on server (NOT firewall), disable MTU discovery on server (expected LAN downgrade 10-20%)

Permanent fix: find the router that drops the ICMP MUST FRAGMENT messages, and convince the person responsible for it to fix the configuration.

Next step: molim da se mreze ukljuce u rjesavanje ovog problema jer je generalno obuhvacen veci broj korisnika.
Moguci problemi kod ADSL korisnika: SMTP, GRE tuneli, pristup web stranicama na ADSL-u iza 2 NAT-a, PPTP itd.
Kao privremeno rjesenje korisnicima preporucivati tuneliranje (SSH, SMTPS, IPSec bez GRE-a, koristenje L2TP-a umjesto PPTP-a)
koje ce "ispeglati" fragmentaciju

Provjereno:

- driver.Global.Prevent_ICMP_MTU_Limit konfigurisano ispravno na SGS-u, ICMP Need Frag message attack nije detektovan
- veliki broj paketa prema ADSL korisnicima ima bad checksume, npr:
cdp.checksum_bad==1 || edp.checksum_bad==1 || ip.checksum_bad==1 || tcp.checksum_bad==1 || udp.checksum_bad==1

- full capture dostupan na oba MTA u /tmp/snooper
```

#60 - 09.10.2008 20:32 - Ernad Husremović

- % završeno promijenjeno iz 50 u 60

bilješka: neven se žali na timeout pri postiranju na redmine, ja evo nisam imao nikakvih problema, a konektovao sam se wirelessom iz hotela holywood.

e sad ko je ovima provajder pojma nemam, garant nije bihnet :)

#61 - 10.10.2008 10:28 - Jasmin Beganović

re: note-53

već sam zatvorio postavke kada sam primio email, ali moram reći da ne očekujem da će ovo uticati na stvar naime ja sam forsiranjem mtu-a 1000 na strani pošiljaoca došao do toga da je komunikacija mimo tunela proradila

to je upravo ovdje fino objašnjeno

```
What happens is that some sending mail servers, mostly UNIX machines, use IP
path MTU discovery with the IP DON'T FRAGMENT bit set. What that translates
to is that they send the data part of the email message in packets that are
just as large as they would across their LAN. When the first large packet
reaches a router that decides the packet is too big for it's network, the
router sends an ICMP MUST FRAGMENT message back to the sending server, which
is fine, because the sending server would respond to that by sending smaller
packets and everything would be fine. The problem is that the ICMP MUST
FRAGMENT messages never make it back to the sending server, therefore the
```

```
message times out. A router in the path between the router that sends the
ICMP MUST FRAGMENT message and the sending server is dropping the ICMP MUST
FRAGMENT feedback messages in a mistaken attempt to protect against certain
attacks. We need to find the router that is dropping these ICMP MUST
FRAGMENT feedback message and have the administrator fix the configuration!
```

smanjivanjem mtu-a na 1000 hernad je zaobišao problematičan router

#62 - 10.10.2008 10:31 - Jasmin Beganović

re note-60 ja sa telekabela nemam nikada problema a isto tako iz office_ze sa bihnetovog adsl-a nismo nikada imali timeout-e

#63 - 13.10.2008 14:24 - Jasmin Beganović

Neven me nije kontaktirao niti je bilo email komunikacije

#64 - 13.10.2008 14:58 - Ernad Husremović

kontaktiraj ga emailom

#65 - 17.10.2008 12:28 - Jasmin Beganović

Konataktirao Nevena

```
Pozdrav Nevene,

Ima li ikakvih pomaka u vezi ovog problema ? Od BIHNET-a nismo dobili nikakve informacije, niti dali su pokuša
li riješiti problem niti jeli riješen.

Molio bih za info.

LP
```

#66 - 28.10.2008 15:11 - Ernad Husremović

Moja kolegica Lejla sa fakulteta je odgovorila, a ja joj uzvratio

lborovina@bhtelecom.ba

```
Hehe,
Selam alejkum Lejla.
Dva puta sam Dženani govorio govorio neću ženu da bihuzurim radi ovoga, nek ide redovnim tokom,
al' eto došo je red na tebe :)

Izgleda da kod vas ništa ne dolazi na pravu adresu dok vam korisnik ne svisne ...
Nebitno sada.

Vezano za donju korespondenciju, mi smo problem time-out-a uočili na lokaciji malta
rama-glas sarajevo (ramaglas@bih.net.ba)

U logovima email servera koji (je) direktno primao poštu (internet SMTP klijenti => adsl.rama-glas.com)
počeli smo dobijati masu timeout error-a pri isporuci pošte od strane nekih servera.
Takav je slučaj bio sa bihnetovim serverima. Rezultat je da je pošta sa recimo @gmail-a uredno dolazila, a sa
@bih.net.ba nikako.

Pošto mi imamo dva internet hosta vanjska, uradili smo sljedeće:
- oborili smo port 25 na adsl.rama-glas.com
- usmjerili da pošta ide na taj vanjski server (ns-2.out.ba)
- napravili SSH tunel ns-2.out.ba => adsl.rama-glas.com
- pošta je bez jednog timeout errora počela dolaziti

Takvo je stanje i sada za rama-glas.com. Znači pošta ide preko ssh tunela.

Ono što me je navelo da reagujem kako sam reagovao je činjenica da smo juče primjetili da neki servisi
(chat server, email smtp server, http server ako se radi o većim stranicama) prijavljuju timeout error-e ili s
e ne mogu pingati.
Prvo smo mislili da se radi o problemima kod tog provajdera (kablovska telecabel zenički) ali smo onda vidjeli
 istu manifestaciju
na našem smtp serveru (to je adsl account hsamrae@bih.net.ba - naša kancelarija u sarajevu gdje ja sjedim) ist
e probleme kao i
kod ramaglas-a. Masu timeout error-a kod pokušaja isporuke.

Ja sam upravo u procesu hitnog prebacivanja email domene na google jer, halali, ali fakat poludismo sa vama.
```

Međutim ostaju mi drugi naši korisnici i naši drugi servisi/serveri kao što je npr. http koji se hostiraju direktno na vašem adsl-u (naravno iza adsl router-a).

Možda je najbolje da pogledaš sve što smo mi radili po ovom pitanju na:

http://redmine.bring.out.ba/issues/show/15409

username/password je:
bhtelecom/xxxx

E sad jedino ne znam da li ćeš dobiti timeout kod pokušaja pristupa radi ovih problema koje mi imamo :(jer se redmine.bring.out.ba hostira na lokaciji
adsl-a hsamrae.

P.S. Poselami Nihada,.... <itd privatni dio poruke>

----- "lejla borovina" <lborovina@bhtelecom.ba> wrote:

> Ernade,
>
> Evo malo prije su nam kolege sa Help deska proslijedile tvoj mail.
>
> Posto je cijela korespondencija malo cudna, nekompletna i na zalost do
> nas
> je stigla tek sada, molim te da nam se javis na telefon ili mailom i
> detaljnije opises problem.
>
> Naime bitno je sta su to male, a sta velike testne poruke i koje to
> veze ima
> sa MPLS MTU-om ili uopste sa MTU-om, sa kojih to lokacija prolaze a sa
> kojih
> ne, na kojim je adresama server Ramaglasa, da li na ADSL liniji ili
> je
> negdje hostiran (prema nasim informacijama Ramaglas je korisnik 9G
> paketa od
> 01.04.2008), da li je u pitanju razmjena mailova izmedju Ramaglasovog
> i
> Bihnet maila, odnosno da li saobracaj prolazi preko nasih SGS-ova ili
> i
> ispred Ramaglasovog servera stoji SGS kojeg odrzava QSS,....
>
> Ukratko, za rjesenje problema jednostavno nemamo dovoljno informacija
> iz
> prilozene prepiske.
>
> Svejedno izuzetno mi je zao sto smo stigli u ovu situaciju, a ona je
> prije
> svega posljedica cinjenice da u rjesavanje problema nisu ukljuceni oni
> koji
> su trebali biti.
>
> Kako god, javi se na moj mobilni ili uradi odgovor na ovu poruku.
>
> Cujemo se.
>
> Lejla Borovina

#67 - 28.10.2008 15:12 - Ernad Husremović

nevjerovatno ali istinito, njima nikada nije prijavljen ovaj problem, a oni su glavni za ove stvari.

#68 - 28.10.2008 15:19 - Ernad Husremović

inače lejla me je zvala i prijavila da ne može preko njihove interne adrese pristupiti

nakon što sam na router-u stavio

root@router-wan-sa-1:/mnt/1/etc# ./routes.sh

```
ip route add 195.222.38.213/32 dev ppp0
```

može, kod nas je sve iz opsega

```
add_route "195.222.0.0/16" "10.0.0.1" "ppp1"
```

freezona osim servera za koje znamo da su public kao što su bihnet serveri, ekonomski fakultet sa.

```
add_route "195.222.33.151" "link" "ppp0"
add_route "195.222.33.152" "link" "ppp0"
add_route "195.222.43.0/24" "link" "ppp0"
```

lejla mi je rekla dati spisak freezone servera da mogu precizno definisati route

#69 - 28.10.2008 15:20 - Ernad Husremović

inače pročitajte i ovaj vezani ticket karakterističnog naziva :) #15316 - bihnetova crna rupa

#70 - 28.10.2008 15:22 - Ernad Husremović

dobra je vijest da ovi problemi nisu rezultat neke smislene akcije tipa - adsl "G" paket - ništa od serverskog pristupa, što smo se mi pobojali

#71 - 28.10.2008 17:04 - Ernad Husremović

testiranje officesa.sigma-com.net

kako sada testirati slanje mail-a

u ns.out.ba sam dodao /var/named/sigma-com.net.zone

```
officesa IN A 89.146.186.231
         MX 10 officesa  <<<<<<<<<<<<<<<<<<<<<
```

na ns-2.out.ba rekao:

```
rndc retransfer officesa.sigma-com.net
```

dig -t mx officesa.sigma-com.net treba dati => MX 10 officesa.sigma-com.net

takođe sam na mail-gw-10, ns-2.out.ba (vps.bring.out.ba) itd otvorio prolaz za test@officesa.sigma-com.net

šaljem ovako (7.5 MB velika poruka)

```
root@bring:~# mail test@officesa.sigma-com.net -s "duga poruka mtu 1460" < test_2.txt
```

ili ovako

```
root@bring:~# mail test@officesa.sigma-com.net -s "duga poruka mtu 3.5 MB" < test.txt
```

idem na zimbra.sigma-com.net logiram se kao test@officesa.sigma-com.net / s..... i vidim je li pošta stigla

#72 - 28.10.2008 17:05 - Ernad Husremović

rezime:

- mogu slati na test@officesa.sigma-com.net
- nakon svakog refresh-a ip-a ovo se mora dodati u ns.out.ba i ns-2.out.ba kako je maloprije objašnjeno

#73 - 28.10.2008 17:09 - Ernad Husremović

lejla mi je rekla da na njihovom router-u stoji 1452 MTU pa je preporučila da taj MTU stavim

komandom na router-u

```
#ifconfig ppo mtu 1452
```

sam to podesio, i to radi.

Međutim, vratio sam kasnije na 1460 MTU i sa tim bez problema prolaze poruke i od 3.5 i od 7.5 MB

izgleda da samo na rama-glas router-u problem možemo reproducirati

#74 - 28.10.2008 17:46 - Ernad Husremović

testiranje rama-glas

da bi testirali rama-glas uradimo sljedeće:

1. uveo sam test@adsl.rama-glas.com adresu na zimbra server rama-glas
2. otvorio port 25 na routeru
3. /var/named/rama-glas.com.zone

```
$TTL 60
@       IN SOA     @    root (
                        0810101430
                        60
                        20
                        3W12h
                        10 )
  MX     10 adsl.rama-glas.com.
;  MX     50 mail-50.sigma-com.net. <<<<<<<<<<<<<<<< ukino
  IN NS     ns.out.ba.
  IN NS     ns-2.out.ba.
adsl IN A 92.36.168.82
  MX     10 adsl.rama-glas.com.  <<<<<<<<<<<<<<<<<<  adsl.rama-glas.com domenu handlira adsl.rama-glas.c
om server
mail CNAME adsl
zimbra CNAME mail
jabber CNAME adsl
web CNAME web-1.sigma-com.net.
www CNAME web
. CNAME adsl
```

4. ovo se naravno mora uraditi nakon svakog refresh_ip-a

test

```
[root@ernadh ~]# mail test@adsl.rama-glas.com -s "duga test" < test.txt
```

#75 - 28.10.2008 17:48 - Ernad Husremović

već sa 2 MB porukom dobijam poznatu grešku

[root@ernadh ~]# mailq

```
-Queue ID- --Size-- ----Arrival Time---- -Sender/Recipient-------
4DEA1230043  2160453 Tue Oct 28 12:41:25  root@smtp.sigma-com.net
(conversation with adsl.rama-glas.com[92.36.168.82] timed out while sending message body)
                                         test@adsl.rama-glas.com

-- 2109 Kbytes in 1 Request.
```

#76 - 28.10.2008 17:49 - Ernad Husremović

pokušaćemo

```
root@router-wan-rg-2:/etc/config# ifconfig ppp0 mtu 1452
root@router-wan-rg-2:/etc/config# ifconfig ppp1 mtu 1452
```

#77 - 28.10.2008 17:50 - Ernad Husremović

napomena, ovaj test se vrši sa ns.out.ba (ne sa ns-2.out.ba jer je na njemu podešen tunel i to ne želim da diram !:

root@ernadh ~# mailq

```
-Queue ID- --Size-- ----Arrival Time---- -Sender/Recipient-------
4DEA1230043* 2160453 Tue Oct 28 12:41:25  root@smtp.sigma-com.net
                                         test@adsl.rama-glas.com

-- 2109 Kbytes in 1 Request.
```

#78 - 28.10.2008 17:51 - Ernad Husremović

kratka poruka na test@adsl.rama-glas.com je prošla bez problema, ali duga zaglavljuje, bez obzira na mtu 1452

#79 - 28.10.2008 18:18 - Ernad Husremović

testiranje rama-glas / 2

i hano-bih.zone takođe treba isto podesiti

```
$TTL 60
@               IN SOA          @       root (
                                0809231338      ; serial
                                60              ; refresh
                                30             ; retry - update retry 15 minutes
                                3W12h           ; expiry - 3 weeks + 12 hours
                                2h )            ; minimum - 2 hours
;                MX              10 adsl.rama-glas.com.   <<<<<<<<<<<<<<< komentarisano
                MX              10 mail-50.sigma-com.net. <<<<<<<<<<<<<<< sve ide na naš ns-2.out.ba
                IN NS           ns.out.ba.
                IN NS           ns-2.out.ba.
                TXT             "hano-bih sarajevo"

web             CNAME           web-1.sigma-com.net.
www             CNAME           web
```

na ns-2.out.ba

```
root@bring:~# rndc retransfer hano-bih.com
```

#80 - 29.10.2008 08:36 - Ernad Husremović

to lborovina at bhtelecom:

```
Selam Lejla,

(1) test ramaglas (ramaglas@bih.net.ba)

sa ovom adresom test@adsl.rama-glas.com možemo reprodukovati problem slanja mail-a na ramaglas@bih.net.ba

http://redmine.bring.out.ba/issues/show/15409#note-74

Kako na ticketu stoji, već sa porukom od 2-3 MB postoje problemi u transferu.

(2) test sigma-com (hsamrae@bih.net.ba)
za testiranje slanja na sigma-com (officesa.sigma-com.net) možemo direktno slati na  adresu test@officesa.sigm
a-com.net
http://redmine.bring.out.ba/issues/show/15409#note-71

Međutim, ovdje je pošta normalno prolazila prilikom slanja poruka od 6-7 MB.

Znači najbolje da se fokusirate na test-ove (1), pa kada skontate šta je razlog tom problemu vjerovatno je to
na čitavoj mreži tako.
```

#81 - 30.10.2008 11:43 - Ernad Husremović

vezano za komentar note-78 i saznanja koje sam pročitao na internetu - long HTTP POST

uradio sljedeće

router ramaglas:

```
root@router-wan-rg-2:/etc/config# ifconfig ppp0 mtu 1452
root@router-wan-rg-2:/etc/config# ifconfig ppp1 mtu 1452
```

zimbra ramaglas

```
root@zimbra:~# ifconfig venet0:0 mtu 1452
```

test from ns.out.ba (128-177-28-71.ip.openhosting.com[128.177.28.71]):

```
[root@ernadh ~]# mail test@adsl.rama-glas.com -s "duga test_2" < test.txt
```

hej poruke su prošle (dvije od 2 MB i jedna od 4 MB) !!

#82 - 30.10.2008 11:47 - Ernad Husremović

test 8MB

root@ernadh ~# cat test_2.txt test_2.txt > test_4.txt

root@ernadh ~# ls -l test_4.txt

```
-rw-r--r--  1 root root 8520568 Oct 30 11:43 test_4.txt
```

root@ernadh ~# mail test@adsl.rama-glas.com -s "duga test_4" < test_4.txt

root@ernadh ~# mailq

```
-Queue ID- --Size-- ----Arrival Time---- -Sender/Recipient-------
A769123003B* 8640887 Thu Oct 30 06:43:38  root@smtp.sigma-com.net
                                         test@adsl.rama-glas.com

-- 8438 Kbytes in 1 Request.
```

root@ernadh ~# mailq

```
Mail queue is empty
```

jupiiiiii ovo izgleda fercera

#83 - 30.10.2008 11:49 - Ernad Husremović

otvaram adsl.ramaglas na MX zapisima

Oct 30 11:40:14 zimbra postfix/smtpd[29136]: timeout after DATA from 128-177-28-71.ip.openhosting.com[128.177.28.71]
root@zimbra:~# tail /var/log/mail.log --lines=10000 | grep "after DATA"

#84 - 30.10.2008 11:52 - Ernad Husremović

```
@       IN SOA     @   root (
                       0810101431
                       60
                       20
                       3W12h
                       10 )
   MX     10 adsl.rama-glas.com.
   MX     50 mail-50.sigma-com.net.
```

ns ns-2.out.ba uradio.

```
root@bring:~# rndc retransfer rama-glas.com
root@bring:~# rndc retransfer hano-bih.com
```

#85 - 30.10.2008 11:53 - Ernad Husremović

bilježim kad mi se desio zadnji timeout

root@zimbra:~# tail /var/log/mail.log --lines=10000 | grep "after DATA"

```
Oct 30 11:40:14 zimbra postfix/smtpd[29136]: timeout after DATA from 128-177-28-71.ip.openhosting.com[128.177.
28.71]
```

#86 - 30.10.2008 11:54 - Ernad Husremović

provjerićemo za par sati ima li problema ili pošta prolazi

#87 - 30.10.2008 13:00 - Ernad Husremović

za sada je sve super

root@zimbra:~# tail /var/log/mail.log --lines=10000 | grep "timeout after DATA"

```
Oct 30 11:40:14 zimbra postfix/smtpd[29136]: timeout after DATA from 128-177-28-71.ip.openhosting.com[128.177.
```

```
28.71]
```

#88 - 30.10.2008 15:33 - Ernad Husremović

i dalje nijednog novog timeout after DATA, pa ovo radi izgleda

#89 - 31.10.2008 09:38 - Jasmin Beganović

```
(09:30:53) hernad: bjasko
(09:31:02) bjasko: reci
(09:31:09) hernad: ja sam sinoć transport mail-50 skinuo sa tunela
(09:31:12) hernad: ide sve direktno
(09:31:16) bjasko: da vidio sam
(09:31:22) hernad: i primarni MX je adsl.rama-glas.com
(09:31:38) hernad: na onom MTU ticketu vidiš
(09:31:43) bjasko: da čitao sam
(09:32:22) bjasko: nešto ne kontam što je sada proferceralo
(09:32:29) hernad: iskreno ni ja
(09:32:35) hernad: novi momenat
(09:32:45) hernad: je što sam na krajnju destinaciju paketa
(09:32:46) bjasko: znam da ste testirali baš ovo isto prije i nije radilo
(09:32:51) hernad: stavio mtu isti ko na routeru
(09:33:10) hernad: jesi provjerio jutros timeout-e ?
(09:33:31) bjasko: ima samo jedan
(09:33:34) bjasko: to je normala
(09:33:55) hernad: ma da to sam i ja sinoć vidio
(09:34:12) bjasko: sve ok fercera
(09:34:18) bjasko: pratriću par dana da vidimo
(09:34:33) bjasko: moguće da su i oni nešto odradili na svom dijelu
09:35
(09:35:34) hernad: stavi ti update svaki put kad pogledaš
(09:35:47) bjasko: ok
(09:35:53) hernad: ja sam te mtu-ove spustio i kod nas na redmine-u
(09:36:02) hernad: izgleda da je to ista stvar i tu bila
(09:36:11) hernad: oni koji su na jačim internet konekcijama
(09:36:32) hernad: i koji imaju kod sebe MTU 1500 su imali probleme sa pristupom našim stranicama
(09:36:39) hernad: recimo postiranjem podataka
(09:36:47) hernad: znaš da je onaj neven pominjao isto timeout
(09:36:55) bjasko: hm ja imam kod kuće taj MTU i nisam imao probleme
(09:37:17) bjasko: čak sam i testirao i MTU fragmentacija radi kada pingam nas ili ramaglas
(09:38:02) bjasko: ma neka samo radi pa šta je god
```

#90 - 31.10.2008 09:39 - Jasmin Beganović

sinoć i jutros pregledao nema timeouta osim jednog što je normala

#91 - 31.10.2008 09:49 - Jasmin Beganović

```
(09:39:09) hernad: pa to očigledno nije pravilo negdje se manifestuje negdje ne
09:40
(09:42:01) bjasko: da to znam ali ako meni ramaglasov router kada ga testiram iz zenice javi da moram fragment
irat paket ispod 1472 neznam zašto bi to drugima uskratio  tako da i dalje mislim da je problem do njihove opr
eme bio
(09:43:12) hernad: to si pingom testirao ?
(09:44:09) bjasko: da iz XP-a
(09:44:19) bjasko: iam ping tu opciju
(09:44:29) hernad: la linux-u nema :) ?
(09:44:55) hernad: man ping
(09:44:59) hernad: -s packetsize
09:45
(09:45:25) bjasko: sekunda da vidim
(09:45:40) bjasko: ima -M
(09:45:47) bjasko: -M hint
              Select Path MTU Discovery strategy.  hint may be either do (prohibit fragmentation, even local o
ne), want (do PMTU dis-
              covery, fragment locally when packet size is large), or dont (do not set DF flag).

(09:46:03) hernad: i šta kaže kada uputiš request > od onog što može "pojesti"
(09:46:26) hernad: de pingaj nas iz zenice da utvrdiš kolika je moguća veličina paketa
(09:47:54) bjasko: ok
(09:48:28) bjasko: heh
```

#92 - 31.10.2008 09:50 - Jasmin Beganović

ovo je dobar rouet npr MTU 1470 mi odmah moj router kaže da je MTU prevelik

bjasko@n-book-bjasko-1:~/admin/backup.rmlh.ba/gitosis-admin$ ping -M do -s 1470 officesa.bring.out.ba

```
PING officesa.sigma-com.net (89.146.136.220) 1470(1498) bytes of data.
From 192.168.44.1 icmp_seq=1 Frag needed and DF set (mtu = 1492)
From 192.168.44.2 icmp_seq=2 Frag needed and DF set (mtu = 1492)
From 192.168.44.2 icmp_seq=2 Frag needed and DF set (mtu = 1492)
```

#93 - 31.10.2008 09:53 - Jasmin Beganović

ovo je interesentno kada sam smanjio na 1455 javio mi se bihnetov router i tu je MTU 1452

bjasko@n-book-bjasko-1:~/admin/backup.rmlh.ba/gitosis-admin$ ping -M do -s 1455 officesa.bring.out.ba

```
PING officesa.sigma-com.net (89.146.136.220) 1455(1483) bytes of data.
From dlp-225.as53.mo.bih.net.ba (195.222.60.225) icmp_seq=1 Frag needed and DF set (mtu = 1452)
From 192.168.44.2 icmp_seq=2 Frag needed and DF set (mtu = 1452)
From 192.168.44.2 icmp_seq=2 Frag needed and DF set (mtu = 1452)
From 192.168.44.2 icmp_seq=2 Frag needed and DF set (mtu = 1452)
```

MTU koji prolazi je 1424 sve preko fragmentiše

bjasko@n-book-bjasko-1:~/admin/backup.rmlh.ba/gitosis-admin$ ping -M do -s 1424 officesa.bring.out.ba
PING officesa.sigma-com.net (89.146.136.220) 1424(1452) bytes of data.

```
1432 bytes from SE400.PPPoE-2268.sa.bih.net.ba (89.146.136.220): icmp_seq=1 ttl=62 time=215 ms
1432 bytes from SE400.PPPoE-2268.sa.bih.net.ba (89.146.136.220): icmp_seq=2 ttl=62 time=175 ms
```

#94 - 31.10.2008 09:55 - Jasmin Beganović

za ramaglas je ta vrijednost sada 1464

bjasko@n-book-bjasko-1:~/admin/backup.rmlh.ba/gitosis-admin$ ping -M do -s 1464 adsl.rama-glas.com

```
PING adsl.rama-glas.com (92.36.168.82) 1464(1492) bytes of data.
1472 bytes from c10k-sa.pppoe10322.sa.bih.net.ba (92.36.168.82): icmp_seq=1 ttl=60 time=222 ms
1472 bytes from c10k-sa.pppoe10322.sa.bih.net.ba (92.36.168.82): icmp_seq=2 ttl=60 time=239 ms
```

#95 - 31.10.2008 09:56 - Jasmin Beganović

izgleda da to kod svih adsl korisnika drugačije ??

#96 - 31.10.2008 09:58 - Ernad Husremović

ma nije jasko, ja sam postavio i na rama-glas i kod nas mtu 1452 juče : note 81

#97 - 31.10.2008 10:00 - Ernad Husremović

ali ja nisam restartovao router-a pa je izgleda došlo do refresh-a

root@router-wan-rg-2:~# ifconfig ppp0

```
ppp0      Link encap:Point-to-Point Protocol
          inet addr:92.36.168.82  P-t-P:92.36.128.1  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1000<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<
          RX packets:131949 errors:0 dropped:0 overruns:0 frame:0
          TX packets:115131 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:3
          RX bytes:51246700 (48.8 MiB)  TX bytes:12464805 (11.8 MiB)
```

#98 - 31.10.2008 10:01 - Ernad Husremović

```
(10:00:35) hernad: vidi je li keno blizu
(10:00:41) hernad: pa restartuj router rama-glas
(10:01:05) hernad: on kod svakog refresh-a ip-a setuje mtu sa kojim je boot-an
(10:01:22) hernad: a ja radi onih problema zaglavljivanja kod restarta to nisam na ramaglas router-u smio učin
iti
(10:01:31) hernad: pa ako reko zaglavi da ga keno trzne ručno
```

#99 - 31.10.2008 10:02 - Ernad Husremović

- % završeno promijenjeno iz 60 u 80

```
(10:02:01) hernad: i hej
(10:02:04) hernad: daj stavi
(10:02:10) hernad: firewall
(10:02:19) hernad: 25 otvori i na /mnt/1/etc/*fw
(10:02:32) hernad: jer sam ja samo u /tmp otvorio
```

#100 - 31.10.2008 10:07 - Ernad Husremović

naš router je ok njega sam nakon promjene config-a restartovao
root@router-wan-sa-1:~# ifconfig ppp0

```
ppp0      Link encap:Point-to-Point Protocol
          inet addr:89.146.136.220  P-t-P:89.146.128.1  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1452  <<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<
          RX packets:327328 errors:0 dropped:0 overruns:0 frame:0
          TX packets:304756 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:3
          RX bytes:68674342 (65.4 MiB)  TX bytes:89449064 (85.3 MiB)
```

#101 - 31.10.2008 10:09 - Ernad Husremović

root@router-wan-rg-2:~# ifconfig ppp0

```
ppp0      Link encap:Point-to-Point Protocol
          inet addr:92.36.168.82  P-t-P:92.36.128.1  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1000  Metric:1
          RX packets:131949 errors:0 dropped:0 overruns:0 frame:0
          TX packets:115131 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:3
          RX bytes:51246700 (48.8 MiB)  TX bytes:12464805 (11.8 MiB)
```

na ramaglasu sam uradio ručno ponov

```
root@router-wan-rg-2:~# ifconfig ppp0  mtu 1452
```

ali i dalje ping kaže da može pojesti 1464

#102 - 31.10.2008 10:10 - Ernad Husremović

```
(09:53:33) hernad: glavno pitanje je ustvari šta se dešava kada se forwarduje paket krajnjoj destinaciji
                npr email serveru da li je on taj koji je mjerodavan da kaže mtu kojim može da se odazove
(09:53:38) hernad: to mi nelogično
(09:54:05) hernad: logično da na ruti router sa najmanjim mtu-om diktira mtu saobraćaja
(09:54:56) hernad: ukratko od nas je stanje bilo client mtu 1500 => adsl router bihnet 1492 => mtu recimo 1000

                   na router-ramaglas => zimbra.rama-glas mtu 1500
(09:55:15) hernad: šta bi trebao klijent dobiti kao info ?
(09:55:24) hernad: po meni mtu najmanje propusnog uređaja
(09:55:38) hernad: ali to izgleda nekada ne funkcioniše
(09:57:48) bjasko: da router bi trebao biti taj koji diktira MTU
```

#103 - 31.10.2008 10:21 - Ernad Husremović

na openwrt forumima http://forum.openwrt.org/viewtopic.php?id=5662

```
...
MTU = 1492
MTU is optimized for PPoE DSL broadband. If not, consider raising MTU to 1500 for optimal throughput.

MSS = 1452
MSS is optimized for PPPoE DSL broadband. If not, consider raising your MTU value.
...
```

#104 - 31.10.2008 10:21 - Ernad Husremović

root@router-wan-rg-2:~# cat /tmp/*fw | grep mss

```
$IPTABLES -A FORWARD -p tcp -m tcp --tcp-flags SYN,RST SYN -j TCPMSS --clamp-mss-to-pmtu
```

#105 - 31.10.2008 10:24 - Ernad Husremović

http://lists.netfilter.org/pipermail/netfilter-devel/2003-October/012941.html

Ideally, the **TCPMSS --clamp-mss-to-pmtu** option would clamp the MSS to the PMTU from the dst to the src. Currently, it uses the PMTU from the packet filter to the dst, which will just be the MTU of the outgoing interface if PMTU discovery hasn't occurred yet. This patch better approximates the full PMTU by including the MTU of the incoming interface ...

#106 - 31.10.2008 10:25 - Ernad Husremović

http://www.fiaif.net/pipermail/fiaif/2004q4/001063.html

To enable ipsec packets through our DSL connection we have **to force a lower MTU on packets forwarded from the LAN to the Internet**.

```
"man iptables" recommends:

    iptables -A FORWARD -p tcp --tcp-flags SYN,RST SYN -j TCPMSS --clamp-mss-to-pmtu
```

#107 - 31.10.2008 10:30 - Ernad Husremović

https://blue-labs.org/howto/mtu-mss.php

Packets with an MTU of 1500 bytes (max transmission unit, or in english, biggest size a packet is allowed to be) cannot pass an interface on a machine where the MTU is less than 1500 bytes. In addition, routers along the way are not allowed to split the packet up and pass it on if the **DF flag is set**. The **DF flag** means Don't Fragment and is set on for almost all TCP traffic.

It breaks the **ability of a remote user to fetch content from your site**. I.e. web pages that are larger than 1500 bytes will simply stall and timeout (Most webpages on the internet are larger than 1500 bytes).

PMTU breaks when naive network administrators set up a **firewall to block all that evil ICMP stuff that evil h4cker dudes use**. This in itself isn't bad ~~if~~ they stripped the DF flag off their packets, but they don't. So what happens now? Packets hit a machine like your PPP dialup, PPPoE tunnel, VPN, or other form of tunneled communications that can't achieve a 1500 byte MTU. When they hit, because the interface can't pass a packet of that size, it drops the packet (remember, the DF flag is set so it isn't allowed to fragment the packet) and issues an ICMP unreachable message with a submessage of "fragmentation needed". Depending on the configuration and location of this particular interface, it may issue a network unreachable or host unreachable, etc.

Now the **ICMP message travels back to the 1500 byte issuer, unfortunately the machine that created the packet will never see that it's packets are being dropped due to the 1500 byte size because their firewall is dropping this ICMP message**.

#108 - 31.10.2008 10:37 - Ernad Husremović

On my firewall, First I set the MTU of my interface to 1492 (or a lower value as applicable, i.e. 1400) instead of 1500, then I **clamp the outbound packets using iptables**. Here is the rule: iptables -A FORWARD -p tcp --tcp-flags SYN,RST SYN -m tcpmss --mss 1400:1536 -j TCPMSS --clamp-mss-to-pmtu. Why do I have that **1400:1536** verb? Well, I only want to modify packets that have MSS values between 1400 and 1536 bytes. Why? Because of my other solution that I use on my internet server. On that machine I have routes set for each of the faulty hosts/networks that I have discovered that have K-mart engineers (my apologies to K-mart Corporation) running their networks. Speaking of that, here's how I handle it. Here is an example route: ip r a 200.199.201.30 via 10.0.0.1 dev eth0 mtu 552 advmss 1. Yes, I very intentionally set the MSS of this to 1 byte of data per packet to severely penalize these networks.

Examples of how to modify the advertised MSS:

For the firewall:

- iptables -A FORWARD -p tcp --tcp-flags SYN,RST SYN -m tcpmss --mss 1400:1536 -j TCPMSS --clamp-mss-to-pmtu

This will calculate the proper size of the MSS based on the MTU of the packet. It is only applied to packets that are traveling through the FORWARD chain and have an original MSS within the 1400 to 1536 range.

For the internal server:

```
ip route add 200.199.201.30 via 10.0.0.1 dev eth0 mtu 552 advmss 1
```

This generates a route which is more specific than the default route. It's design is to set a fixed MSS of 1 for this particular host. That means that there is the TCP header and only 1 byte for the payload. So to send "hello", it would take five packets (one packet for each letter) [thanks David DeSimone for the correction].

are you naive network admin ?

You simply need to allow ICMP unreachable -- fragmentation needed messages (type 3, code 4). Your visitors ~~and~~ internal users will love you for it. You should never blindly just block or allow packets willy nilly.

Methods of doing this:

Ipchains - Allow ICMP type 3, code 4, then block all other type 3

```
    ipchains -A output -s <$external_FW_interface_IP> 3 -d 0.0.0.0/0 4 -p ICMP -j ACCEPT
    ipchains -A output -s <$internal_network_CIDR> 3 -d 0.0.0.0/0 4 -p ICMP -j ACCEPT
    ipchains -A output -s <$external_FW_interface_IP> 3 -p ICMP -j DENY
    ipchains -A output -s <$internal_network_CIDR> 3 -p ICMP -j DENY
```

Iptables - Allow only ICMP type 3, code 4 to be passed through in the FORWARD table, drop everything else. Remember, the FORWARD chain has no effect on the packets destined for this firewall, only packets traveling through it. For packets destined for this firewall you add the same rules to the INPUT chain.

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

#109 - 31.10.2008 10:38 - Ernad Husremović

man iptables

```
       [!] --mss value[:value]
              Match a given TCP MSS value or range.
```

#110 - 31.10.2008 10:39 - Ernad Husremović

man iptables

```
       --clamp-mss-to-pmtu
              Automatically clamp MSS value to (path_MTU - 40).
```

#111 - 31.10.2008 10:39 - Ernad Husremović

hernad@nmraka-1:~$ dict clamp

```
1 definition found

From English-Croatian Freedict Dictionary [fd-eng-cro]:

  clamp

  pričvrstiti alat, spajati, spojnica, spona, uklještenje
```

#112 - 31.10.2008 11:00 - Ernad Husremović

testovi pinga ns.out.ba => officesa.

[root@ernadh ~]# ping -M do -s 1452 officesa.sigma-com.net

```
From ernadh.user.openhosting.com (128.177.28.71) icmp_seq=0 Frag needed and DF set (mtu = 1452)
```

#113 - 31.10.2008 11:01 - Ernad Husremović

aha 1452-28 je vrijednost koju naš host može vratiti

[root@ernadh ~]# ping -M do -s 1424 officesa.sigma-com.net

```
PING officesa.sigma-com.net (89.146.150.136) 1424(1452) bytes of data.
1432 bytes from SE400.PPPoE-5768.sa.bih.net.ba (89.146.150.136): icmp_seq=0 ttl=50 time=166 ms
```

#114 - 31.10.2008 11:02 - Ernad Husremović

samo jedan bajt više i dobijamo poruku kako treba.

Sigma-com se ponaša ispravno

[root@ernadh ~]# ping -M do -s 1425 officesa.sigma-com.net

```
PING officesa.sigma-com.net (89.146.150.136) 1425(1453) bytes of data.
From SE400.PPPoE-1.sa.bih.net.ba (89.146.128.1) icmp_seq=0 Frag needed and DF set (mtu = 1452)
```

#115 - 31.10.2008 11:07 - Ernad Husremović

što se tiče podešenja firewall-a, **isto se ponaša** (javlja ispravan mut) i u stanju bez da je sav icmp saobraćaj omogućen

staro stanje
root@router-wan-sa-1:~# iptables -L | grep icmp

```
In_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
RULE_14    icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
RULE_14    icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
RULE_14    icmp --  anywhere             anywhere            icmp type 0 code 0
RULE_14    icmp --  anywhere             anywhere            icmp type 8 code 0
```

novo stanje

root@router-wan-sa-1:~# iptables -L | grep icmp

```
In_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp any state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
In_RULE_9  icmp --  anywhere             anywhere            icmp any state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp any state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
Out_RULE_9  icmp --  anywhere             anywhere            icmp any state NEW
RULE_14    icmp --  anywhere             anywhere            icmp type 0 code 0 state NEW
RULE_14    icmp --  anywhere             anywhere            icmp type 8 code 0 state NEW
RULE_14    icmp --  anywhere             anywhere            icmp type 0 code 0
RULE_14    icmp --  anywhere             anywhere            icmp type 8 code 0
```

#116 - 31.10.2008 11:09 - Ernad Husremović

testovi rama-glas

rama-glas ping ašićare **laže** !

```
[root@ernadh ~]# ping -M do -s 1425 adsl.rama-glas.com
PING adsl.rama-glas.com (92.36.169.22) 1425(1453) bytes of data.
1433 bytes from 92.36.169.22: icmp_seq=0 ttl=50 time=185 ms
1433 bytes from 92.36.169.22: icmp_seq=1 ttl=50 time=189 ms

--- adsl.rama-glas.com ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 185.315/187.355/189.396/2.085 ms, pipe 2
[root@ernadh ~]# ping -M do -s 1424 adsl.rama-glas.com
PING adsl.rama-glas.com (92.36.169.22) 1424(1452) bytes of data.
1432 bytes from 92.36.169.22: icmp_seq=0 ttl=50 time=178 ms

--- adsl.rama-glas.com ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 178.802/178.802/178.802/0.000 ms, pipe 2
[root@ernadh ~]# ping -M do -s 1460 adsl.rama-glas.com
PING adsl.rama-glas.com (92.36.169.22) 1460(1488) bytes of data.
1468 bytes from 92.36.169.22: icmp_seq=0 ttl=50 time=189 ms

--- adsl.rama-glas.com ping statistics ---
2 packets transmitted, 1 received, 50% packet loss, time 1004ms
rtt min/avg/max/mdev = 189.148/189.148/189.148/0.000 ms, pipe 2
[root@ernadh ~]# ping -M do -s 1462 adsl.rama-glas.com
```

```
PING adsl.rama-glas.com (92.36.169.22) 1462(1490) bytes of data.
1470 bytes from 92.36.169.22: icmp_seq=0 ttl=50 time=189 ms
1470 bytes from 92.36.169.22: icmp_seq=1 ttl=50 time=181 ms

--- adsl.rama-glas.com ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1002ms
rtt min/avg/max/mdev = 181.872/185.754/189.636/3.882 ms, pipe 2
```

ovdje sve prolazi !

#117 - 31.10.2008 11:12 - Ernad Husremović

za velike pakete rama-glas **ništa ne uzvraća** !

[root@ernadh ~]# ping -M do -s 1472 adsl.rama-glas.com

```
PING adsl.rama-glas.com (92.36.169.22) 1472(1500) bytes of data.
```

samo se zakoči

isto i sada

[root@ernadh ~]# ping -M do -s 1471 adsl.rama-glas.com

```
PING adsl.rama-glas.com (92.36.169.22) 1471(1499) bytes of data.
```

#118 - 31.10.2008 11:14 - Ernad Husremović

koji je limit

[root@ernadh ~]# ping -M do -s 1463 adsl.rama-glas.com

```
PING adsl.rama-glas.com (92.36.169.22) 1463(1491) bytes of data.
1471 bytes from 92.36.169.22: icmp_seq=0 ttl=50 time=190 ms
```

[root@ernadh ~]# ping -M do -s 1464 adsl.rama-glas.com

```
PING adsl.rama-glas.com (92.36.169.22) 1464(1492) bytes of data.
1472 bytes from 92.36.169.22: icmp_seq=0 ttl=50 time=189 ms
```

na **1465** (a to je +28 = 1493) ping počne zaglavljivati

[root@ernadh ~]# ping -M do -s 1465 adsl.rama-glas.com

```
PING adsl.rama-glas.com (92.36.169.22) 1465(1493) bytes of data.
```

znači ping tih velikih paketa prema ramaglas se ne javlja on jednostano počne zaglavljivati !!!!!!!!!!! očigledno na routeru ispred našeg routera - znači na bihnetovom router-u

#119 - 31.10.2008 14:45 - Ernad Husremović

adsl.rama-glas.com još testova

kako sam maloprije rekao zaglavljuje višim vrijednostima

root@ernadh ~# ping -M do -s 1472 adsl.rama-glas.com

```
PING adsl.rama-glas.com (92.36.169.22) 1472(1500) bytes of data.
```

a onda kada paket **prevali 1500 dobijamo odziv !**

root@ernadh ~# ping -M do -s 1472 adsl.rama-glas.com

```
From ernadh.user.openhosting.com (128.177.28.71) icmp_seq=0 Frag needed and DF set (mtu = 1500)
```

znači kada pokušamo poslati 1501 byte javi nam se **neko** i kaže da je njegov mtu 1500 znači javi nam se bihnetov adsl router, ko drugi

#120 - 31.10.2008 14:46 - Ernad Husremović

```
----- "lejla borovina" <lborovina@bhtelecom.ba> wrote:

> Ernade,
>
>
> Sto se tice MTU-a kojeg ramaglas prijavljuje 'cudno' moguce je i da
> adsl modem koji je ispred routera (ukoliko je tako, odnosno ukoliko ne
> koristite router sa adsl linijom) unosi to laganje, odnosno i na modemima
> postoji
> podesenje MTU-a, pa bi mogli provjetiti ako vam se jos da, da li su u
> pitanju isti modemi i koliki je MTU satovan na njima.
>
da može biti modem, ali on je za nas blackbox, vi ste ga dostavili,
mislim da mu mi ni ne možemo prići (pogledati/promjeniti config). Ako možemo,
proslijedi instrukcije

> Sto se adrese tice podeseni su na diskonekt na 12 sati, a sto se moze i
> vidjeti sa samog uredjaja (inate jos 10 sati i 51 minutu za ovu
> konekciju). Desava se da ponekad nakon diskonekta korisnik ponovo dobije istu IP
> adresu koju je imao i prije, jer ovaj Cisco pametnjakovic po nekom svom
> internom
> algoritmu pridruzuje slobodne adrese iz poola i to izgleda po FOFI
> principu,
> pa ukoliko se nakon diskonekta odmah pokrene konekt i ukoliko te ne
> pretekne neki drugi zahrjev adresa se moze ponoviti.
> Culi smo za ovo i nista mu nemozemo (a bismo da znamo kako).
>
Nama je svejedno :)

> Hajde da onda ramaglas-mail problem ostavimo do iduce sedmice pa da
> onda mogu razduziti smetnju na help-desku i objasniti im sta da preporuce
> korisnicima u slicnim slucajevima.
ok
```

#121 - 31.10.2008 14:48 - Ernad Husremović

vezano za adsl modem ?

```
(31.10.2008 13:44:13) hernad: vezano za ramaglas adsl
(31.10.2008 13:44:13) bjasko: reci
(31.10.2008 13:44:18) hernad: taj njihov adsl modem
(31.10.2008 13:44:35) bjasko: ja
(31.10.2008 13:44:37) hernad: lejla pominje i njega da li mi njemu možemo prići ?
(31.10.2008 13:44:52) hernad: navodno da na njemu nije nešto loše itd
(31.10.2008 13:44:52) bjasko: nemožemo on je u bridge-u
(31.10.2008 13:45:00) hernad: a to znači šta ?
(31.10.2008 13:45:01) bjasko: nema ništa on samo propušta
(31.10.2008 13:45:08) bjasko: router ostvaruje konekciju
(31.10.2008 13:45:18) hernad: ne dira/mjenja mtu itd
(31.10.2008 13:45:36) bjasko: on ja u stanju u kakvom ga je pošta isporučila i kada je u bridg-e ništa se nemo
že mjenjati
```

#122 - 31.10.2008 14:52 - Ernad Husremović

ubuntu forums [SOLVED] change mtu

#123 - 31.10.2008 14:53 - Jasmin Beganović

re: **znači kada pokušamo poslati 1501 byte javi nam se neko i kaže da je njegov mtu 1500 znači javi nam se bihnetov adsl router, ko drugi**

javlja ti se tvoj host 128.177.28.71 da mu je MTU prevelik (1500) je njegov limit

sam MTU lan-a je 1500 znači iznad ovoga bi ti morala odmah tvoja mrežna vratiti DF

#124 - 31.10.2008 15:35 - Ernad Husremović

pametan si jale brate, pametan

#125 - 31.10.2008 15:36 - Ernad Husremović

iptables TCPMSS target

The TCPMSS target can be used to alter the MSS (Maximum Segment Size) value of TCP SYN packets that the firewall sees. The MSS value is used to control the maximum size of packets for specific connections. Under normal circumstances, this means the size of the **MTU (Maximum Transfer Unit) value, minus 40 bytes**. This is used to overcome some **ISP's and servers that block ICMP fragmentation needed packets**, which can result in really weird problems which can mainly be described such that everything works perfectly from your firewall/router, **but your local hosts behind the firewall can't exchange large packets**. This could mean such things as mail servers being able to send small mails, but not large ones, web browsers that connect but then hang with no data received, and ssh connecting properly, but scp hangs after the initial handshake. In other words, everything that uses any large packets will be unable to work.

The TCPMSS target is able to solve these problems, by changing the size of the packets **going out through a connection**. Please note that **we only need to set the MSS on the SYN packet** since the hosts take care of the MSS after that. The target takes two arguments.

#126 - 31.10.2008 15:36 - Ernad Husremović

http://en.wikipedia.org/wiki/Maximum_transmission_unit

#127 - 31.10.2008 15:38 - Ernad Husremović

- Fajl mss-talk.pdf dodano

MSS - maximum segment size

#128 - 31.10.2008 15:41 - Ernad Husremović

Path MTU discovery (wikipedia)

The Internet Protocol defines the "path MTU" of an Internet transmission path as the **smallest MTU of any of the IP hops of the "path" between a source and destination**. Put another way, the path MTU is the largest packet size that traverse this path without suffering fragmentation.

RFC 1191 describes "Path MTU discovery", a technique for determining the path MTU between two IP hosts. It works by **setting the DF (Don't Fragment) option in the IP headers of outgoing packets**. **Any device along the path whose MTU is smaller than the packet will drop such packets and send back an ICMP "Destination Unreachable (Datagram Too Big)" message containing its MTU**, allowing the source host to reduce its assumed path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.

Unfortunately, **increasing numbers of networks drop ICMP traffic (e.g. to prevent denial-of-service attacks)**, which prevents path MTU discovery from working. One often detects such blocking in the cases where a connection works for low-volume data but hangs as soon as a host sends a large block of data at a time. For example, with IRC a connecting client might see up to the ping message, but get no response after that. This is because the large set of welcome messages are sent out in packets bigger than the real MTU. Also, in an IP network, the path from the source address to the destination address often gets modified dynamically, in response to various events (load-balancing, congestion, outages, etc.) - this could result in the path MTU changing (sometimes repeatedly) during a transmission, which may introduce further packet drops before the host finds the new safe MTU.

Most Ethernet LANs use an **MTU of 1500 bytes** (modern LANs can use **Jumbo frames, allowing for an MTU up to 9000 bytes**), however **border protocols like PPPoE will reduce this**. This causes path MTU discovery to come into effect with the possible result of making some sites behind badly-configured firewalls unreachable. One can possibly work around this, depending on which part of the network one controls; for example one can change the MSS (maximum segment size) in the initial packet that sets up the TCP connection at one's firewall.

This problem has surfaced more frequently since the introduction of Windows Vista which introduces the 'Next Generation TCP/IP Stack'. This implements "Receive Window Auto-Tuning that continually determines the optimal receive window size by measuring the bandwidth-delay product and the application retrieve rate, and adjusts the maximum receive window size based on changing network conditions".[2] This has been seen to fail in conjunction with older routers and firewalls that appeared to work with other operating systems. It is most often seen in ADSL routers and can often be rectified by a firmware update.

#129 - 31.10.2008 16:22 - Ernad Husremović

- Fajl iptables_tutorial.zip dodano

http://iptables-tutorial.frozentux.net/iptables-tutorial.html

#130 - 01.11.2008 12:04 - Ernad Husremović

http://www.snailbook.com/faq/mtu-mismatch.auto.html

SSH Frequently Asked Questions
My SSH session hangs part way through logging on, when I generate a lot of output from my shell, try to scp or sftp a file, or attempt to run an X11 application. I have a firewall, NAT or packet filter.
Contributed by Darren Tucker (dtucker at zip.com.au) and edited by RES; previously published here
Short Answer
You probably have an MTU/fragmentation problem. For each network interface on both client and server set the MTU to 576, eg ifconfig eth0 mtu 576. If the problem goes away, read on.
Long Answer
Long answer: At each routing hop, IP packets bigger than the outgoing interface's Maximum Transmission Unit (MTU) get fragmented. Only the first fragment has TCP port numbers. Firewalls often behave badly in the presence of packet fragmentation, dropping everything but the first fragment since the subsequent ones can't be matched against the firewall rules. Some NAT configuration (eg many-to-one NAT or port address translation) can't match the fragments against their translation state tables.

Arguably, such devices should perform packet reassembly first so as to properly consider fragmented packets. However, this is more complicated and so is often not done. Also, this feature would raise a possible starvation attack against the packet filter, by sending many bogus initial fragments and causing the device to store them for reassembly with subsequent packets which will never come.

Logging in and using the shell will normally generate relatively small packets, and so the initial connection proceeds normally ; however if do you something that generates a lot of data (eg cat'ing a big file or starting an X Windows application), you may generate a packet bigger than the MTU.

Let's say it's a 1500 byte IP packet and the router has 2 different MTU's (say 1500 & 1484) and no firewall. When the router goes to forward it, the packet is too big for the interface **MTU (1484), so the router breaks it into 2 fragments, 0 and 1**. **Fragment 0 contains the first 1484 bytes** (including the TCP source and dest ports) and fragment **1 contains the remaining 16 bytes**. Both fragments are sent on to their destinations.

When the **first fragment reaches its target**, it's held **by the IP stack until the remaining fragments arrive**, at which time the IP packet is **reassembled and passed up the stack to TCP**. **If all fragments are not received by the timeout**, the entire IP packet is discarded and an ICMP "timeout during reassembly" error is sent back.

Now add your firewall, which drops fragment 1. Your 1500 byte IP packet times out during reassembly and TCP retries, by sending another 1500 byte packet. Repeat. Eventually, TCP will time out and you'll get a connection termination.

IP stack parameters (such as Path MTU Discovery) and external variables (such as the MTU's of all the hops between hosts) can also affect whether or not a given connection will have this problem.

#131 - 01.11.2008 12:09 - Ernad Husremović

dobar blog na ovu temu

http://packetlife.net/blog/2008/aug/18/path-mtu-discovery/

root@ernadh ~# tracepath redmine.bring.out.ba

```
 1:  ernadh.user.openhosting.com (128.177.28.71)             0.184ms pmtu 1500
 1:  128-177-27-1.ip.openhosting.com (128.177.27.1)          0.963ms
 2:  209.249.9.123 (209.249.9.123)                        asymm  3   0.748ms
 3:  xe-2-0-0.er2.dca2.us.above.net.31.125.64.in-addr.arpa (64.125.31.118)   1.332ms
 4:  64.125.31.210 (64.125.31.210)                        asymm  5   1.773ms
 5:  above-telia.iad10.us.above.net (64.125.13.190)       asymm  6   2.173ms
 6:  ldn-bb2-pos6-0-0.telia.net (213.248.65.209)          asymm  9  75.427ms
 7:  ffm-bb1-link.telia.net (80.91.254.204)               asymm 13  93.528ms
 8:  win-b2-link.telia.net (80.91.252.77)                 asymm 10 103.753ms
 9:  croatiantelecom-ic-124945-win-b2.c.telia.net (213.248.84.190) asymm 12 109.446ms
10:  gdr01-gdr10-1.ip.t-com.hr (195.29.252.189)           asymm 12 124.375ms
11:  195.29.249.70 (195.29.249.70)                        asymm 12 125.812ms
12:  sama6513.bih.net.ba (195.222.32.228)                 asymm 13 120.497ms
13:  uplink-6.se400.bih.net.ba (195.222.42.54)            asymm 14 124.411ms
14:  SE400.PPPoE-1.sa.bih.net.ba (89.146.128.1)           129.387ms pmtu 1452
15:  no reply
16:  no reply
17:  no reply
18:  no reply
...
```

#132 - 01.11.2008 12:12 - Ernad Husremović

tracepath fino dođe do mog router-a, a kada to pokušamo sa ramaglas-om šupica tach zapne na ovom hostu dlp-65.max2.sa-mlt.bih.net.ba (195.222.42.65)

root@ernadh ~# tracepath adsl.rama-glas.com

```
 1:  ernadh.user.openhosting.com (128.177.28.71)             0.143ms pmtu 1500
 1:  128-177-27-1.ip.openhosting.com (128.177.27.1)          0.770ms
 2:  209.249.9.123 (209.249.9.123)                        asymm  3   0.752ms
 3:  xe-2-0-0.er2.dca2.us.above.net.31.125.64.in-addr.arpa (64.125.31.118)   1.409ms
 4:  64.125.31.210 (64.125.31.210)                        asymm  5   2.081ms
 5:  above-telia.iad10.us.above.net (64.125.13.190)       asymm  6   2.347ms
 6:  prs-bb2-link.telia.net (80.91.254.215)               asymm  9  82.300ms
 7:  ffm-bb2-link.telia.net (80.91.248.66)                asymm 13 114.349ms
 8:  win-b2-link.telia.net (80.91.252.77)                 asymm 10 103.510ms
 9:  croatiantelecom-ic-124945-win-b2.c.telia.net (213.248.84.190) asymm 12 116.612ms
10:  gdr01-gdr10-1.ip.t-com.hr (195.29.252.189)           asymm 12 119.656ms
11:  195.29.249.82 (195.29.249.82)                        asymm 12 127.885ms
12:  sama6513.bih.net.ba (195.222.32.228)                 asymm 13 121.890ms
13:  dlp-65.max2.sa-mlt.bih.net.ba (195.222.42.65)        asymm 14 122.087ms    <<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<
<<<<<
14:  no reply
```

```
15:  no reply
16:  no reply
17:  no reply
```

#133 - 01.11.2008 12:17 - Ernad Husremović

na pomenutom blogu fino piše:

When a host needs to transmit data out an interface, it references the interface's Maximum Transmission Unit (MTU) to determine how much data it can put into each packet. Ethernet interfaces, for example, have a default MTU of 1500 bytes, not including the Ethernet header or trailer. This means a host needing to send a TCP data stream would typically use the **first 20** of these 1500 bytes for the **IP header**, the **next 20 for the TCP header**, and as much of the remaining **1460 bytes** as necessary for the **data payload**. Encapsulating data in maximum-size packets like this allows for the least possible consumption of bandwidth by protocol overhead.

#134 - 01.11.2008 12:21 - Ernad Husremović

jedan komentar ovog blog-a:

Disabling or removing the DF bit is never a good idea. Modern routers are not designed to do fragmentation (and it's not even possible to do fragmentation with IPv6). For example the 6500/7600 series punts all fragmented packets to the MSFC and is not handled in hardware by the PFC.

Obviously the best fix is to find out where the ICMP breakdown is occuring and fixing it (thus fixing PMTUD). Another option is **tcp-mss-adjust** (although personally I think that routers shouldn't dip into L4 headers but that just opinion). Removing the DF-bit or never setting it should never be a valid work around!

ovaj tcp-mss-adjust je cisco-v feature

The TCP MSS Adjustment feature enables the configuration of the maximum segment size (MSS) for transient packets that traverse a router, specifically TCP segments in the SYN bit set, when PPP over Ethernet (PPPoE) is being used in the network. PPPoE truncates the Ethernet maximum transmission unit (MTU) 1492, and if the effective MTU on the hosts (PCs) is not changed, the router in between the host and the server can terminate the TCP sessions. The ip tcp adjust-mss command specifies the MSS value on the intermediate router of the SYN packets to avoid truncation.

#135 - 01.11.2008 14:11 - Ernad Husremović

promjena mtu-a se ne vidi izvana ?!

uradio na router-wan-sa-1 ifconfig ppp0 mtu 1492

i to lan moj fino kaže:

hernad@nmraka-1:/data/devel/rails/rest_1$ tracepath ns-2.out.ba

```
 1:  nmraka-1.local (192.168.45.166)                         0.348ms pmtu 1500
 1:  192.168.45.254 (192.168.45.254)                        30.218ms
 1:  192.168.45.254 (192.168.45.254)                         2.236ms
 2:  192.168.45.254 (192.168.45.254)                         2.216ms pmtu 1492
 2:  SE400.PPPoE-1.sa.bih.net.ba (89.146.128.1)             45.160ms
 3:  uplink-3.se400.bih.net.ba (195.222.42.41)              42.505ms
 4:  195.222.32.226 (195.222.32.226)                        44.292ms
 5:  s15-3.border0-ip2.net.telekom.hu (84.1.64.17)          69.418ms asymm 10
 6:  PO9-0.bud-001-access-100.interoute.net (84.233.170.133)  53.087ms asymm 10
 7:  xe-3-0-0-0.bud-001-score-1-re0.interoute.net (84.233.147.117)  96.987ms asymm 22
 8:  ae2-0.prg-001-score-2-re0.interoute.net (84.233.138.213)  97.946ms asymm 21
 9:  ae0-0.prg-001-score-1-re0.interoute.net (84.233.138.205)  99.810ms asymm 20
10:  ae2-0.fra-006-score-2-re0.interoute.net (84.233.138.210) 127.924ms asymm 19
11:  ae0-0.fra-006-score-1-re0.interoute.net (84.233.207.93)  96.445ms asymm 18
12:  ae1-0.ams-koo-score-2-re0.interoute.net (84.233.190.49) 100.318ms asymm 14
13:  ae0-0.ams-koo-score-1-re0.interoute.net (84.233.190.1)  95.218ms asymm 16
14:  ams-ix.ams01.mzima.net (195.69.144.73)                  98.154ms asymm 16
15:  eos4-0.cr01.lhr01.mzima.net (216.193.255.101)          115.988ms asymm 17
16:  eos1-0.cr01.lga02.mzima.net (216.193.255.45)           190.875ms asymm 18
17:  xe1-0.cr01.lga01.mzima.net (216.193.255.213)           185.422ms asymm 19
18:  eos3-2.cr01.ord01.mzima.net (216.193.255.54)           203.610ms asymm 20
19:  eos1-22.cr02.sjc02.mzima.net (216.193.255.154)         248.024ms asymm 21
20:  eos1-23.cr01.sea01.mzima.net (216.193.255.174)         271.992ms asymm 22
21:  xe0-2.cr01.sea02.mzima.net (216.193.255.234)           279.356ms asymm 23
22:  ge1-spry.cust.sea02.mzima.net (72.37.232.34)           247.834ms asymm 25
23:  po1-core0-tuk.wa.spry.com (64.79.223.2)                249.777ms
24:  vps11-031.vpslink.com (66.249.15.76)                   250.730ms asymm 27
25:  209.40.203.155 (209.40.203.155)                        252.764ms reached
     Resume: pmtu 1492 hops 25 back 38
```

ali ne i kontra-upit (internet ns-2.out.ba => officesa)

root@bring:~# tracepath www.bring.out.ba

```
 1:  209.40.203.155 (209.40.203.155)                          0.093ms pmtu 1500
 1:  vpsl1-031.vpslink.com (66.249.15.76)                     0.050ms
 1:  vpsl1-031.vpslink.com (66.249.15.76)                     0.033ms
 2:  po1-br0-tuk.wa.spry.com (64.79.223.1)                    3.990ms asymm  3
 3:  g2.8-br1-tuk.wa.spry.com (64.79.223.26)                  4.070ms asymm  4
 4:  66.162.128.21 (66.162.128.21)                            3.079ms asymm  6
 5:  peer-01-so-0-0-0-0.snjs.twtelecom.net (64.129.248.17)   20.657ms asymm  7
 6:  te8-2.mpd01.sjc01.atlas.cogentco.com (154.54.6.237)     20.957ms asymm  8
 7:  te8-3.ccr02.sfo01.atlas.cogentco.com (154.54.2.137)     23.748ms asymm  9
 8:  te2-3.ccr02.mci01.atlas.cogentco.com (154.54.24.110)    60.903ms asymm 10
 9:  te9-2.ccr01.ord01.atlas.cogentco.com (154.54.25.82)     66.734ms asymm  8
10:  te9-2.mpd01.bos01.atlas.cogentco.com (154.54.7.82)     181.822ms asymm 18
11:  te2-4.mpd02.lon01.atlas.cogentco.com (130.117.0.45)    183.778ms asymm 17
12:  te1-1.ccr01.par01.atlas.cogentco.com (130.117.1.122)   180.987ms asymm 17
13:  te3-1.mpd01.par02.atlas.cogentco.com (130.117.1.229)   177.971ms asymm 16
14:  te4-4.ccr01.fra03.atlas.cogentco.com (130.117.1.66)    179.951ms asymm 15
15:  te1-1.ccr01.vie01.atlas.cogentco.com (130.117.3.22)    199.592ms
16:  te2-2.ccr01.muc01.atlas.cogentco.com (130.117.0.166)   186.042ms asymm 14
17:  te1-1.ccr01.vie01.atlas.cogentco.com (130.117.3.22)    195.972ms asymm 15
18:  149.6.174.30 (149.6.174.30)                            193.030ms
19:  149.6.174.30 (149.6.174.30)                            192.942ms asymm 18
20:  gtr11-gtr10.ip.t-com.hr (195.29.240.98)                195.019ms asymm 21
21:  gdr11-gtr11.ip.t-com.hr (195.29.240.101)               198.849ms asymm 20
22:  195.29.249.82 (195.29.249.82)                          198.997ms asymm 19
23:  195.29.249.70 (195.29.249.70)                          204.923ms asymm 20
24:  sama6513.bih.net.ba (195.222.32.228)                   214.986ms
25:  SE400.PPPoE-1.sa.bih.net.ba (89.146.128.1)             205.965ms pmtu 1452
25:  uplink-4.se400.bih.net.ba (195.222.42.46)              207.883ms asymm 24
26:  no reply
27:  no reply
28:  no reply
29:  no reply
30:  no reply
31:  no reply
     Too many hops: pmtu 1452
     Resume: pmtu 1452
```

u čemu je kvaka - kada izlazimo onda naš router javlja (192.168.45.254) da je na 1492, a kada paket ulazi onda SE400.PPPoE-1.sa.bih.net.ba javlja naš mtu.

#136 - 01.11.2008 18:58 - Ernad Husremović

nakon restarta router-wan-sa-1 dobija se mtu 1492 - izgleda da bihnetova strana peer-a uzima mtu od router-a prilikom uspostavljanja pppoe konekcije i tu info daje drugima

#137 - 10.11.2008 11:30 - Jasmin Beganović

- Prioritet promijenjeno iz Urgentno u Nizak

za sada se više timeouti ne javljaju pa ovo ide na niski prioritet

#138 - 10.11.2008 11:31 - Ernad Husremović

- Status promijenjeno iz Dodijeljeno u Zatvoreno

- % završeno promijenjeno iz 80 u 100

ma zatvaramo ga onda, ticket je ionako prevelik.

Fajlovi

mss-talk.pdf	1,71 MB	31.10.2008	Ernad Husremović
iptables_tutorial.zip	6,75 MB	31.10.2008	Ernad Husremović