
Posted by [maxx](#) on Thu, 29 Oct 2009 05:06:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

<http://forum.nag.ru/forum/index.php?showtopic=52117>.

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

```
iptables -A FORWARD -p tcp --tcp-flags SYN,RST SYN -j TCPMSS --set-mss <value>
```

Posted by [RXL_](#) on Thu, 29 Oct 2009 20:57:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Posted by [maxx](#) on Fri, 30 Oct 2009 10:39:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

SIOCSIFMTU: Invalid argument

Posted by [RXL_](#) on Sat, 31 Oct 2009 22:08:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
# ip l s dev eth1 mtu 1540
# ip l sh dev eth1
3: eth1: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1540 qdisc pfifo_fast qlen 1000
    link/ether 00:30:48:2b:dd:e3 brd ff:ff:ff:ff:ff:ff
# ip l s dev eth1 mtu 1500
# ip l sh dev eth1
3: eth1: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:30:48:2b:dd:e3 brd ff:ff:ff:ff:ff:ff
```

```
# ip l sh dev eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc prio qlen 1000
    link/ether 00:80:48:5d:cb:65 brd ff:ff:ff:ff:ff:ff
# ip l s dev eth0 mtu 1540
SIOCSIFMTU: Invalid argument
```

Posted by [maxx](#) on Tue, 03 Nov 2009 13:20:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Posted by [RXL_](#) on Tue, 03 Nov 2009 18:40:30 GMT

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

```
# ip t a tun1 mode ipip remote 1.2.3.4 dev eth0
# ip l s dev tun1 mtu 1500
# ip l sh dev tun1
6: tun1@eth0: <POINTOPOINT,NOARP> mtu 1500 qdisc noop
   link/ipip 0.0.0.0 peer 1.2.3.4
# ip l sh dev eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast qlen 100
   link/ether 00:17:31:5e:1d:cd brd ff:ff:ff:ff:ff:ff
```

Posted by [maxx](#) on Wed, 04 Nov 2009 13:53:26 GMT

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

- 1500.

Frag needed and DF set (mtu = 1480)

"unreachable - need to frag (mtu 1500)".

13: veth102.0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1540 qdisc noqueue
inet 192.168.11.1/30 scope global veth102.0

7: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1540 qdisc noqueue
inet 192.168.11.2/30 brd 192.168.11.3 scope global eth0

11: tun0@NONE: <POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1520 qdisc noqueue


```
link/ipip 192.168.11.2 peer 192.168.22.150
inet 192.168.100.13 peer 192.168.100.14/30 brd 192.168.100.15 scope global tun0
```

```
[root@vr102 /]# ip ro
192.168.1.0/24 via 192.168.100.14 dev tun0
default via 192.168.11.1 dev eth0
```

```
[root@vr102 /]# ping -s 1452 -M do -I 192.168.2.1 192.168.1.1
```

```
[root@vr102 /]# tcpdump -nv -i tun0
```

```
tcpdump: listening on tun0, link-type RAW (Raw IP), capture size 96 bytes
04:43:08.156747 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: ICMP (1), length: 1480)
192.168.2.1 > 192.168.1.1: ICMP echo request, id 17231, seq 1, length 1460
04:43:08.160526 IP (tos 0x0, ttl 255, id 45198, offset 0, flags [DF], proto: ICMP (1), length: 1480)
192.168.1.1 > 192.168.2.1: ICMP echo reply, id 17231, seq 1, length 1460
...
```

```
[root@vr102 /]# tcpdump -nv -i eth0
```

```
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
04:43:08.156767 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: IPIP (4), length: 1500)
192.168.11.2 > 192.168.22.150: IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: ICMP (1),
length: 1480) 192.168.2.1 > 192.168.1.1: ICMP echo request, id 17231, seq 1, length 1460
04:43:08.160526 IP (tos 0x0, ttl 252, id 45199, offset 0, flags [none], proto: IPIP (4), length: 1500)
192.168.22.150 > 192.168.11.2: IP (tos 0x0, ttl 255, id 45198, offset 0, flags [DF], proto: ICMP (1),
length: 1480) 192.168.1.1 > 192.168.2.1: ICMP echo reply, id 17231, seq 1, length 1460
...
```

veth102.0.

```
[root@vr102 /]# ping -s 1453 -M do -I 192.168.2.1 192.168.1.1
```

```
[root@vr102 /]# tcpdump -nv -i tun0
```

```
04:43:15.777312 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: ICMP (1), length: 1481)
```

```
192.168.2.1 > 192.168.1.1: ICMP echo request, id 17487, seq 1, length 1461
```

```
04:43:16.777686 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: ICMP (1), length: 1481)
```

```
192.168.2.1 > 192.168.1.1: ICMP echo request, id 17487, seq 2, length 1461
```

```
[root@vr102 /]# tcpdump -nv -i eth0
```

```
04:43:15.777321 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: IPIP (4), length: 1501)
```

```
192.168.11.2 > 192.168.22.150: IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: ICMP (1),
```

```
length: 1481) 192.168.2.1 > 192.168.1.1: ICMP echo request, id 17487, seq 1, length 1461
```

```
04:43:15.777354 IP (tos 0xc0, ttl 64, id 62606, offset 0, flags [none], proto: ICMP (1), length: 576)
```

```
192.168.11.1 > 192.168.11.2: ICMP 192.168.22.150 unreachable - need to frag (mtu 1500),  
length 556
```

```
IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: IPIP (4), length: 1501) 192.168.11.2 >
```

```
192.168.22.150: IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto: ICMP (1), length: 1481)
```

```
192.168.2.1 > 192.168.1.1: ICMP echo request, id 17487, seq 1, length 1461[[icmp]
```

192.168.11.1 (HN).

```
04:43:15.777354 IP (tos 0xc0, ttl 64, id 62606, offset 0, flags [none], proto: ICMP (1), length: 576)
```

```
192.168.11.1 > 192.168.11.2: ICMP 192.168.22.150 unreachable - need to frag (mtu 1500),  
length 556
```

From 192.168.2.1 icmp_seq=3 Frag needed and DF set (mtu = 1480)

From 192.168.2.1 icmp_seq=3 Frag needed and DF set (mtu = 1480)

From 192.168.2.1 icmp_seq=3 Frag needed and DF set (mtu = 1480)

...

