NAT Observations Regarding TCP, DCCP, SCTP and ICMP

UNIVERSITY OF HELSINKI

Ability to NAT ICMP messages

ICMP destination port unreachable in response to UDP traffic and Time-to-Live exceeded are the only ICMP messages that all NATs except *nw1* translate.
Two of the tested NATs fail to recalculate all checksums in the ICMP payload, 14 more fail to calculate the transport layer checksum if it exists. One NAT transforms ICMP messages to TCP reset segments but fails to do so properly (bad sequence number).

DNS proxy support

All NATs in our testbed include a DNS proxy but only 13 of the proxies support DNS over TCP.
Deployment of DNSSEC increases the DNS reply sizes beyond what fits in a single UDP packet, which makes TCP more suitable for DNS. Missing support for DNS over TCP may cause problems as DNSSEC deployment progresses.

DCCP and SCTP compatibility

DCCP is clearly entirely unsupported by NAT boxes while SCTP seems to be doing better. It is still unclear whether the NATs actually understand SCTP or use a simplistic NAT algorithm where a single client can use the protocol with a single server outside.

Tag	DCCP: Conn.	DNS over TCP	DNS over UDP	ICMP: Host Unreach.	SCTP: Conn.	TCP: Reass. Time. Ex.	TCP: Frag. Needed	TCP: Param. Prob.	TCP: Source Quench	TCP: Src. Route Fail.	TCP: TTL Exceeded	TCP: Host Unreach.	TCP: Net Unreach.	TCP: Port Unreach.	TCP: Proto. Unreach.	UDP: Reass. Time Ex.	UDP: Frag. Needed	UDP: Param. Prob.	UDP: Source Quench	UDP: Src. Route Fail	UDP: TTL Exceeded	UDP: Host Unreach.	UDP: Net Unreach.	UDP: Port Unreach.	UDP: Proto. Unreach.
al		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ap			●		●	●	●	●	●	●	●	●	●	●	●	●	●		●	●	●	●	●	●	●
as1			●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
be1			●		●	●	●	●	●	●	●	●	●	●	●	●	●		●	●	●	●	●	●	●
be2			●		●	●	●	●	●	●	●	●	●	●	●	●	●		●	●	●	●	●	●	●
bu1		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
dl1			●		●	●	●	●	●	●	●	●	●	●	●	●	●		●	●	●	●	●	●	●
dl10			●			●	●	●	●	●	●	●	●	●	●	●	●		●	●	●	●	●	●	●
dl2	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
dl3		●	●		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
dl4		●									●			●								●			
dl5		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
dl6	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
dl7	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
dl8		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
dl9		●									●			●								●			
ed	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
je	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ls1		●	●		●			●								●	●	●	●	●	●	●	●	●	●
ls2		●	●													●	●		●	●	●	●	●	●	●
ls3		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ls5		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ng1		●	●		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ng2		●	●		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ng3	●	●	●		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ng4	●	●	●		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
ng5		●	●													●				●	●	●	●	●	●
nw1		●																							
owrt	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
smc		●									●			●								●			
te		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
to	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
we	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
zy1		●	●		●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●

Table: Ability to NAT ICMP messages, DNS proxy support, DCCP and SCTP compatibility

TCP Throughput

Many NATs are clearly unable to perform at the link bandwidth (100 Mb/s) and many have even more performance problems with full-duplex traffic.
The clearly suboptimal performance probably goes unnoticed with home users because the Internet connection is still typically slower than the NAT, but this is already changing for the poorest performers.

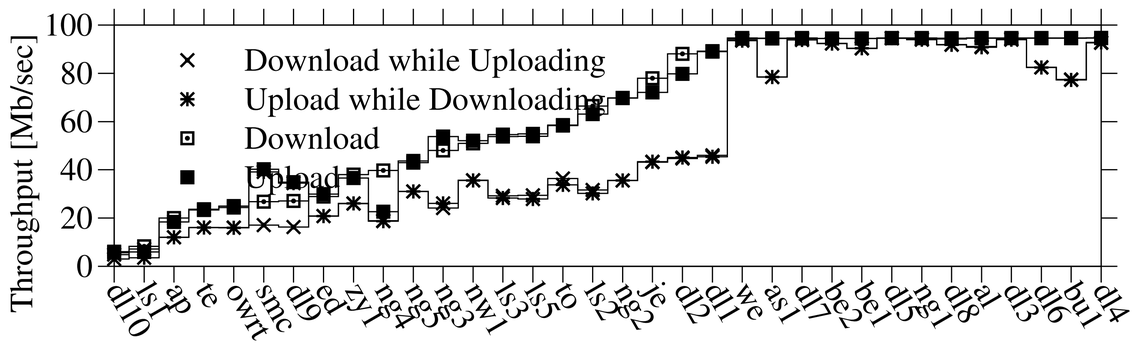

Maximum Number of TCP Bindings

The maximum number of simultaneous TCP bindings varies from 16 in the low end to 1024 in the high. The test is from a single system behind a NAT to a single server:port. Determining the maximum number of bindings from many clients to many servers or ports is still underway.

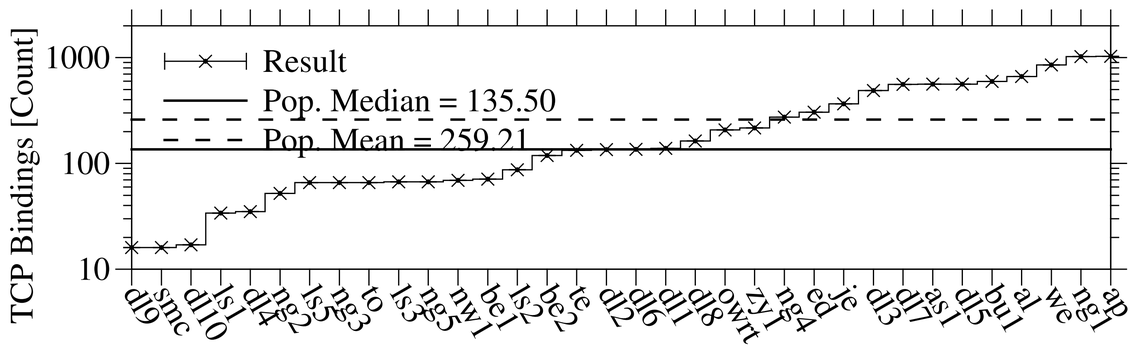

Queuing Delay

Queuing (and processing) delay stays within reasonable bounds for most of the devices, but there is a statistically significant group that build up delays of 50 ms or more.

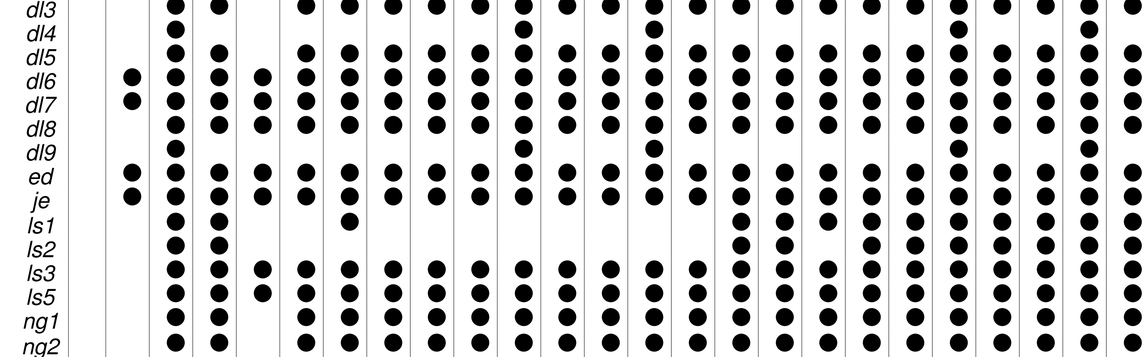

Binding Timeout

Even though it is feasible to make a NAT that never times out TCP connections (contrary to the case for UDP) unless there is excess demand for them, many have a timeout, some even irritatingly short one. About half of the tested NATs have a timeout lower than the IETF recommended 7440 seconds, which is set slightly above the standard TCP keepalive interval of 7200 seconds. Many either had no timeout or the timeout was over 24 hours.

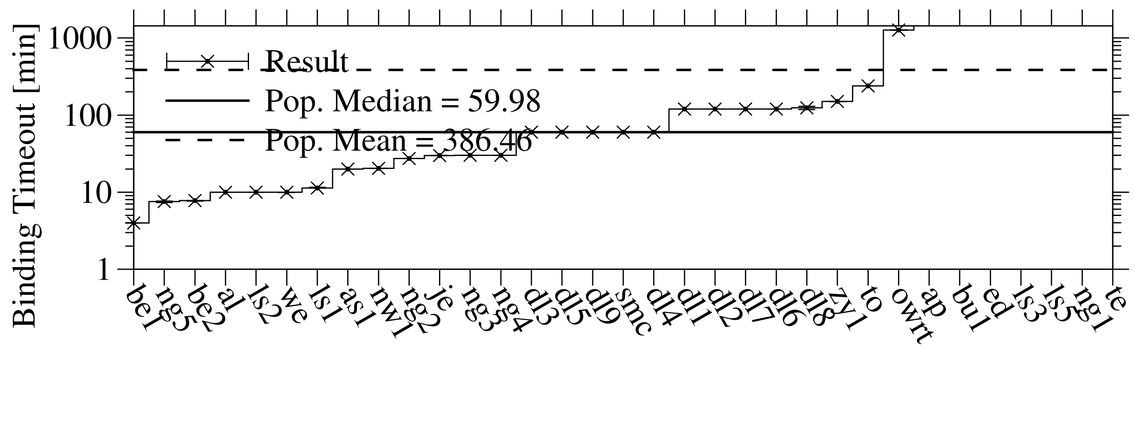

Seppo Hätönen (UH), Aki Nyrhinen (UH), Lars Eggert (NRC), Pasi Sarolahti (HIIT), Stephen Strowes (UG), Markku Kojo (UH)