# Native SCTP, DCCP, UDP-Lite and Home Gateway NATs

**Runa Barik** (UiO), Michael Welzl, Gorry Fairhurst, Thomas Dreibholz, Ahmed Elmokashfi, Stein Gjessing

maprg @104th IETF Meeting
Prague
28$^{th}$ March 2019

# Methodology

- Detecting support for native SCTP, DCCP, UDPLite using transport header checksum (and Vtag for native SCTP).
- Using local testbed to study NAT boxes
  - Off-the-shelf equipment tests
  - *netfilter* in Linux
  - IPF, PF and IPFW in FreeBSD
- Lessons learnt

# Off-the-shelf equipment tests

- Dlink: DIR 655-A2, A3, B1; DIR 619-Ax; DI-614+-B2
- Jensen: Air:Link WBR 7954 v2,v3; AirLink 1000Gv2 (A)
- Linksys: E2500, WRT54G/ GL/GS v1.1, WRT54G, E4200
- Netgear: WGR 614v7, WGR 614v9, WNDR3400
- Topcom: WBR 254G, BR 604
- TP-LINK: TL-MR3020 v1, TL-WR703N
- 3G modem: WR3G050-02 (Spi59-YJ v2.0)
- ZyXEL: P8702N, P-2812HNU-F3
- Edimax: BR-6574N (A)
- Xiaomi: Router 3C

OS and versions:

- Linux 2.4-Linux 4.4,ThreadX, VxWorks, OpenWRT MiWiFi ROM, Unknown.

| Protocol with transport header checksum | Observation | Obs. No. |
|---|---|---|
| UDP with zero checksum | NA(P)Ting with zero checksum | Obs-1 |
| DCCP, SCTP, UDP-Lite and unassigned number | NAT'ing, i.e. no transport header update | Obs-2 |
| DCCP, SCTP, UDP-Lite and unassigned number | No NAT'ing | Obs-3 |
| DCCP, SCTP, UDP-Lite and unassigned number | Dropping | Obs-4 |

Table: Behavior of transport protocols across the middleboxes

- UDP with zero checksum remains intact in all devices.
- VxWorks OS follows Obs-4 for new transport protocols.
- Jensen, Topcom and some Dlink devices follow Obs-3.
- The remaining devices follow IP-level NATing (success for SCTP, but not for DCCP and UDPLite).

# *Netfilter* in Linux

- Linux kernel (version 3.18.109 for MIPS architecture) in TP-Link TL-MR3020, using OpenWRT.
- Netfilter: conntrack_proto_*X* and nat_proto_*X*
- X = {SCTP, DCCP, UDPLite}

conntrack_proto_*X* : Responsible for transport header verification, NATing, state machine. It fails for port collision.

nat_proto_*X* : Responsible for port-mapping and checksum update.

- Linux netfilter supports SCTP, DCCP and UDPLite.
- Netfilter changes the SCTP port on collision.
- *sysctl* variable *nf_conntrack_checksum* verifies the checksum on incoming packets

# FreeBSD

- Used FreeBSD 11.2 in a x86-64 PC.
- Firewall variants of FreeBSD: IPF, PF and IPFW
- No support for DCCP and UDPLite.
- Only IPFW/*libalias* supports SCTP.

# FreeBSD

| **Protocol** | **IPFW** (first / later clients) | **PF** (first / later clients) | **IPF** (first / later clients) |
|---|---|---|---|
| UDP with zero checksum | Obs-1 / Obs-1 | Obs-1 / Obs-1 | Obs-1 / Obs-1 |
| DCCP, UDP-Lite | Obs-2 / Obs-2* | Obs-2 / Obs-4 | Obs-2 / Obs-3 |
| SCTP | Obs-2 / Obs-2 | Obs-2 / Obs-4 | Obs-2 / Obs-3 |

Table: Behavior of transport protocols in FreeBSD NAT firewalls (*: the response packets are forwarded to the last client)

| **Observation** | **Obs. No.** |
|---|---|
| Obs-1 | NA(P)Ting with zero checksum |
| Obs-2 | NAT'ing, i.e. no transport header update |
| Obs-3 | No NAT'ing |
| Obs-4 | Dropping |

Table: Behavior of transport protocols across the middleboxes

# Idle-timeout of NAT devices

Smaller idle-timeout is good for SCTP.

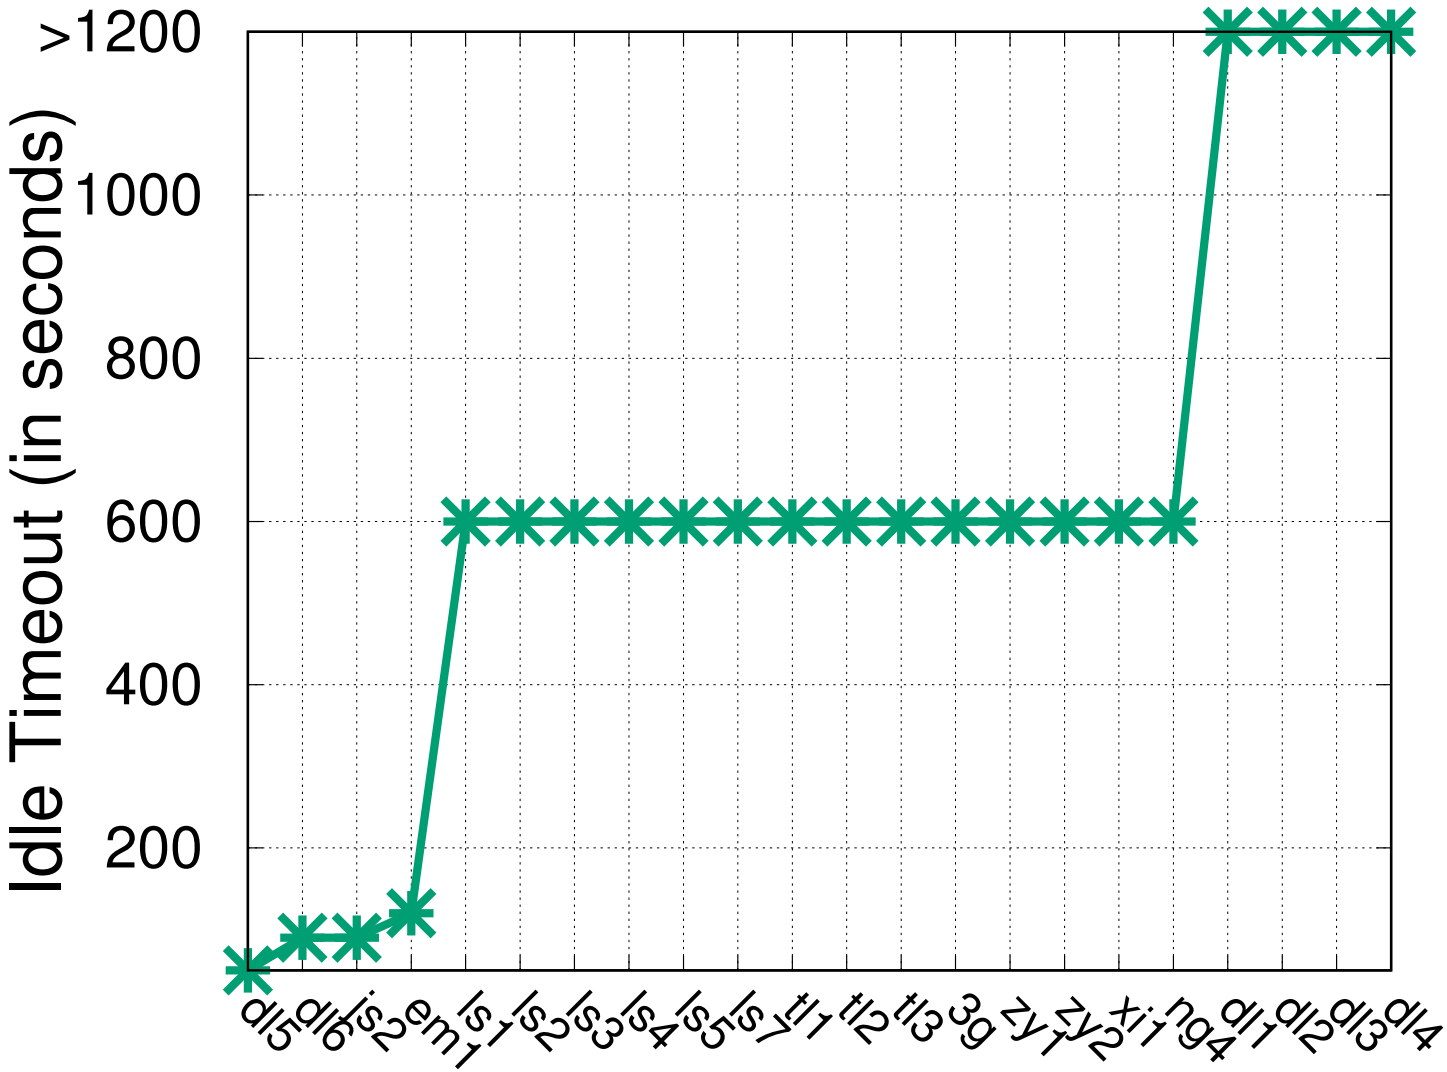

# Lessons learnt

- Configurable support in the device setting for checksum verification or new transport protocols.
- Emulating UDP-Lite with UDP with zero checksum.
- Native protocols are not enabled by default.
- An unusual mechanism or a new option in UDP or TCP has greater chances of success than a new protocol.
- Avoiding the pseudo-header for checksum calculation could improve the chances of NAT traversals.
- To support multi-homing feature, individual connections to a NAT should look like single-homed ones.

# Q&A?

# SCTP Multi-homing with IPFW/*libalias*

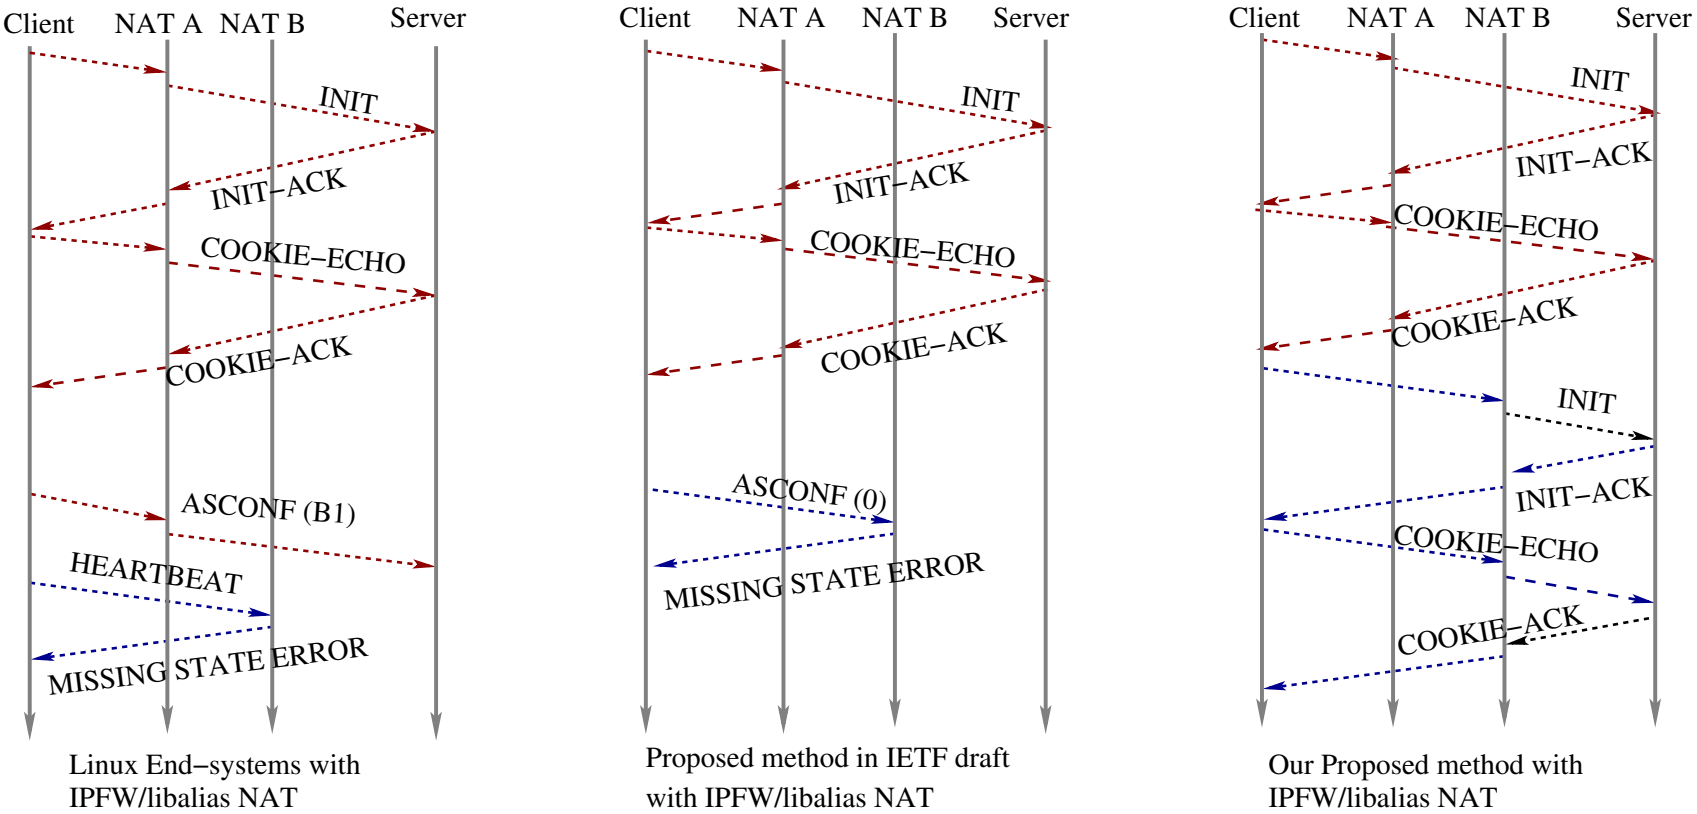


Sequence diagram showing SCTP multi-homing support