IPv6 Extension Headers in the Real World v3.0

Fernando Gont <fgont@si6networks.com>
Jen Linkova <furry@google.com>

IEPG 91
Honolulu, Hawaii, USA. November 9, 2014

“IPv6 EHs in Real World” saga

- Fernando Gont @ IEPG 88:

  > 50% drop rate for small EHs (e.g. DOH of 8 bytes)

  > 40% drop rate for Fragmented traffic

  > 90% drop rate for large EHs (e.g. DOH of 1K)

- Tim Chown & Fernando @ IETF 89:

  > 60% of packet drops >= 7 hops from destination

- Jen Linkova & Fernando @ IETF 90
  - Packets largely dropped at non-destination AS

Some unanswered questions

- What about EHs such as (IPsec's) ESP?
- IPv6 EHs drops for “fun & profit”?

Measurement Results

IPv6 Extension Headers Filtering Measurements with RIPE Atlas

Methodology

To each destination from each probe:

For $PROTOCOL in (“ICMP”, “UDP”):

- control measurement ($PROTOCOL traceroute)
- 9 $PROTOCOL traceroute tests:
    - Hop-by-Hop Options:
        - 8 bytes, 512 bytes, 1024 bytes
    - Destination Options
        - 8 bytes, 512 bytes, 1024 bytes
    - Hop-by-Hop + Destination Options
        - 8 bytes + 8 bytes
        - 128 bytes + 128 bytes
        - 512 bytes + 512 bytes

Testing Topology

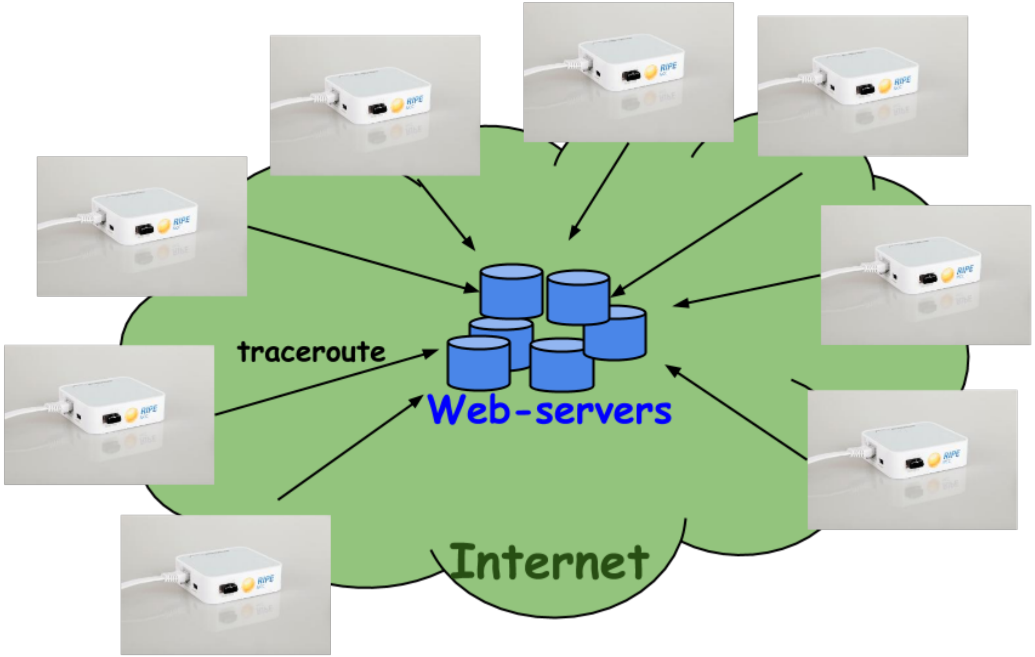

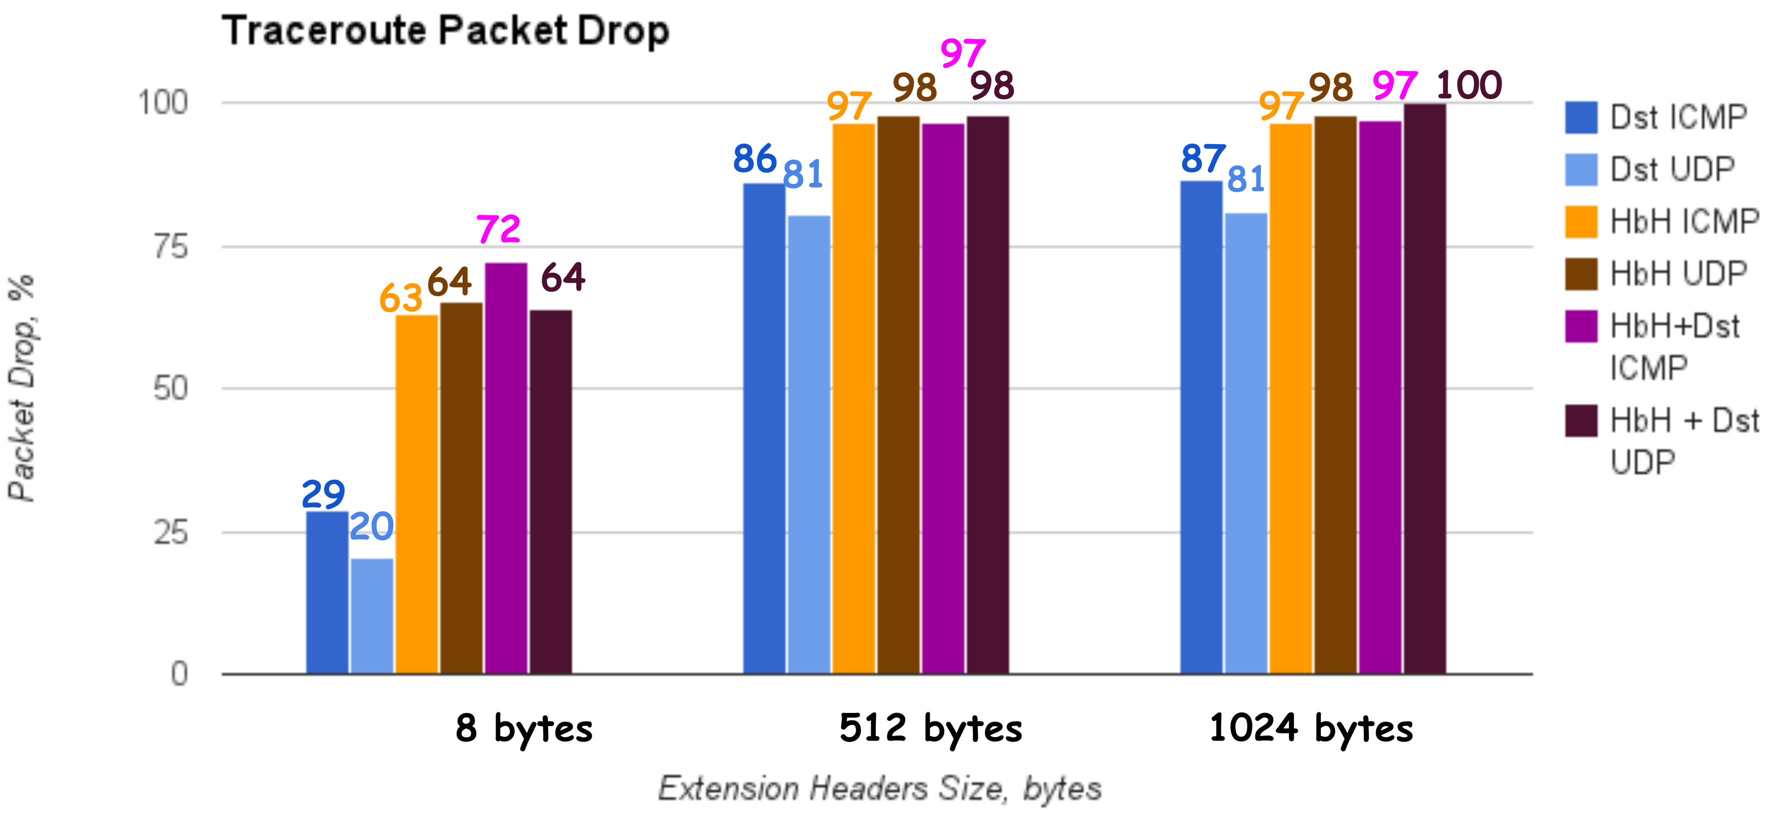
Traceroute Packet Drop
Packet Drop, %
100
75
50
25
0
29
20
63
64
72
64
86
81
97
98
97
98
87
81
97
98
97
100
8 bytes
512 bytes
1024 bytes
Extension Headers Size, bytes
Dst ICMP
Dst UDP
HbH ICMP
HbH UDP
HbH+Dst ICMP
HbH + Dst UDP

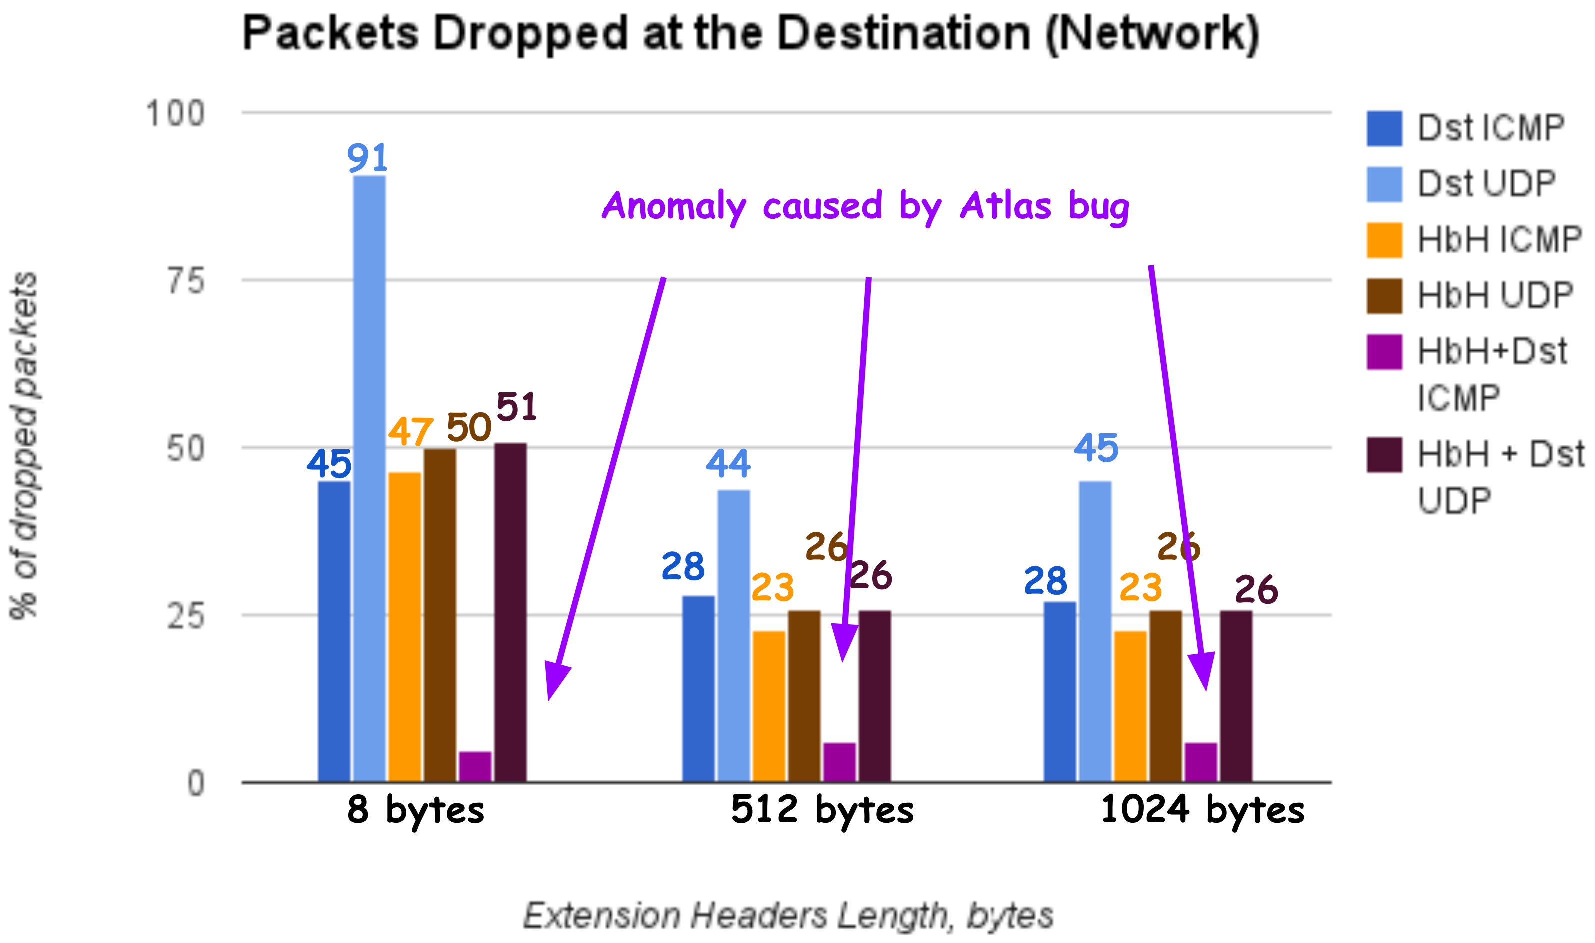

Packets Dropped at the Destination (Network)

Anomaly caused by Atlas bug

Where Are Packets Dropped?

- Finding origin AS for each traceroute hops
- Ignoring invalid IPs/link-local/ULAs/etc
- Comparing 'AS_PATH' for control test and the measurement;
    - If AS_PATH for failed test has length 0 or 1:
        - packet could not leave the origin network
    - If last AS in AS_PATH for failed test is destination AS or PHP AS from the control test:
        - packet was dropped in the destination network or on its edge

Packets Dropped in Transit

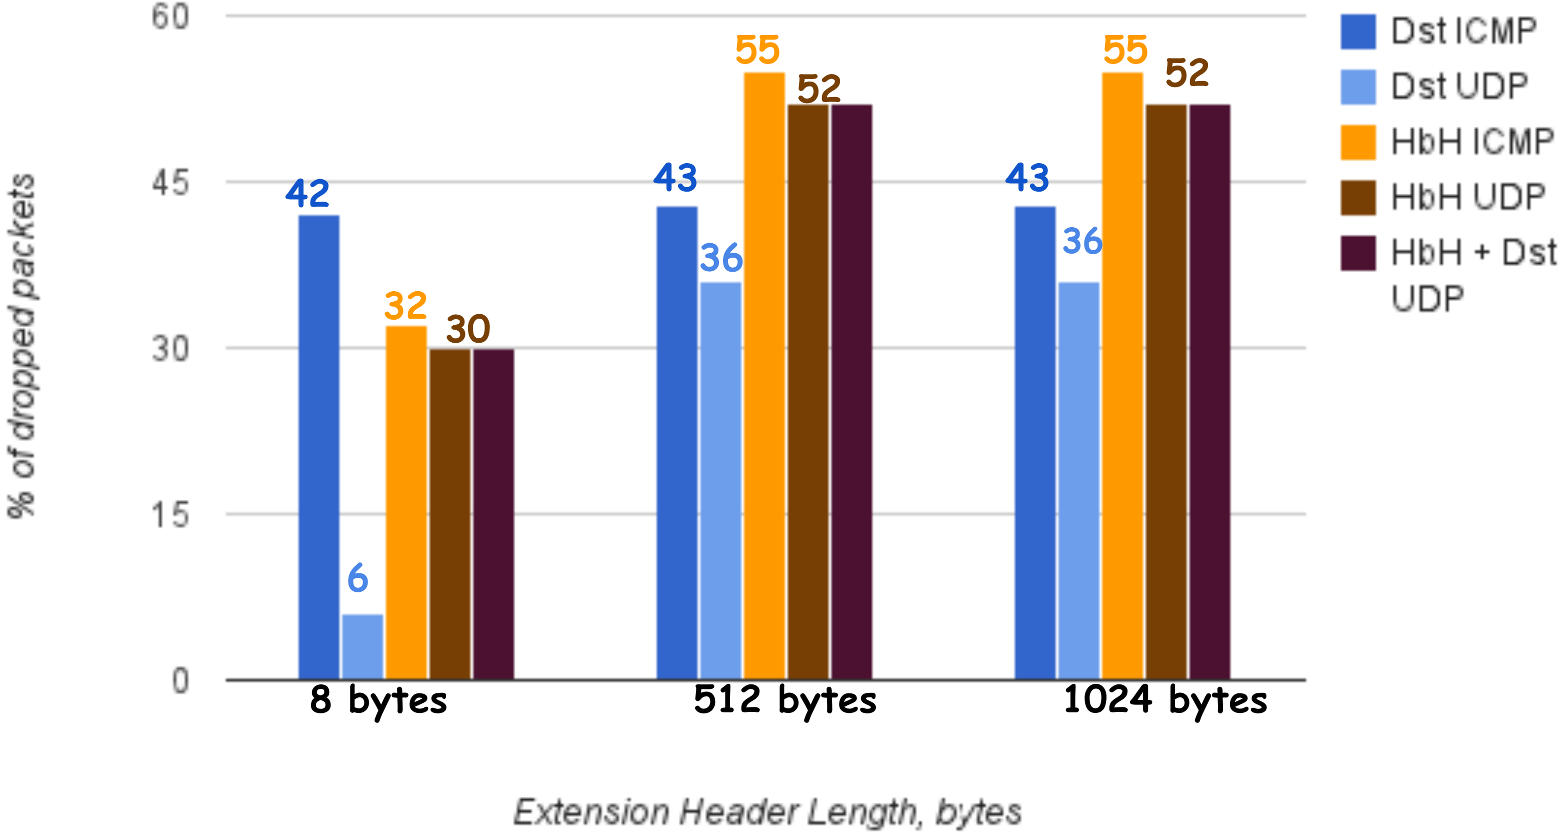

Packets Dropped in the Origin Network

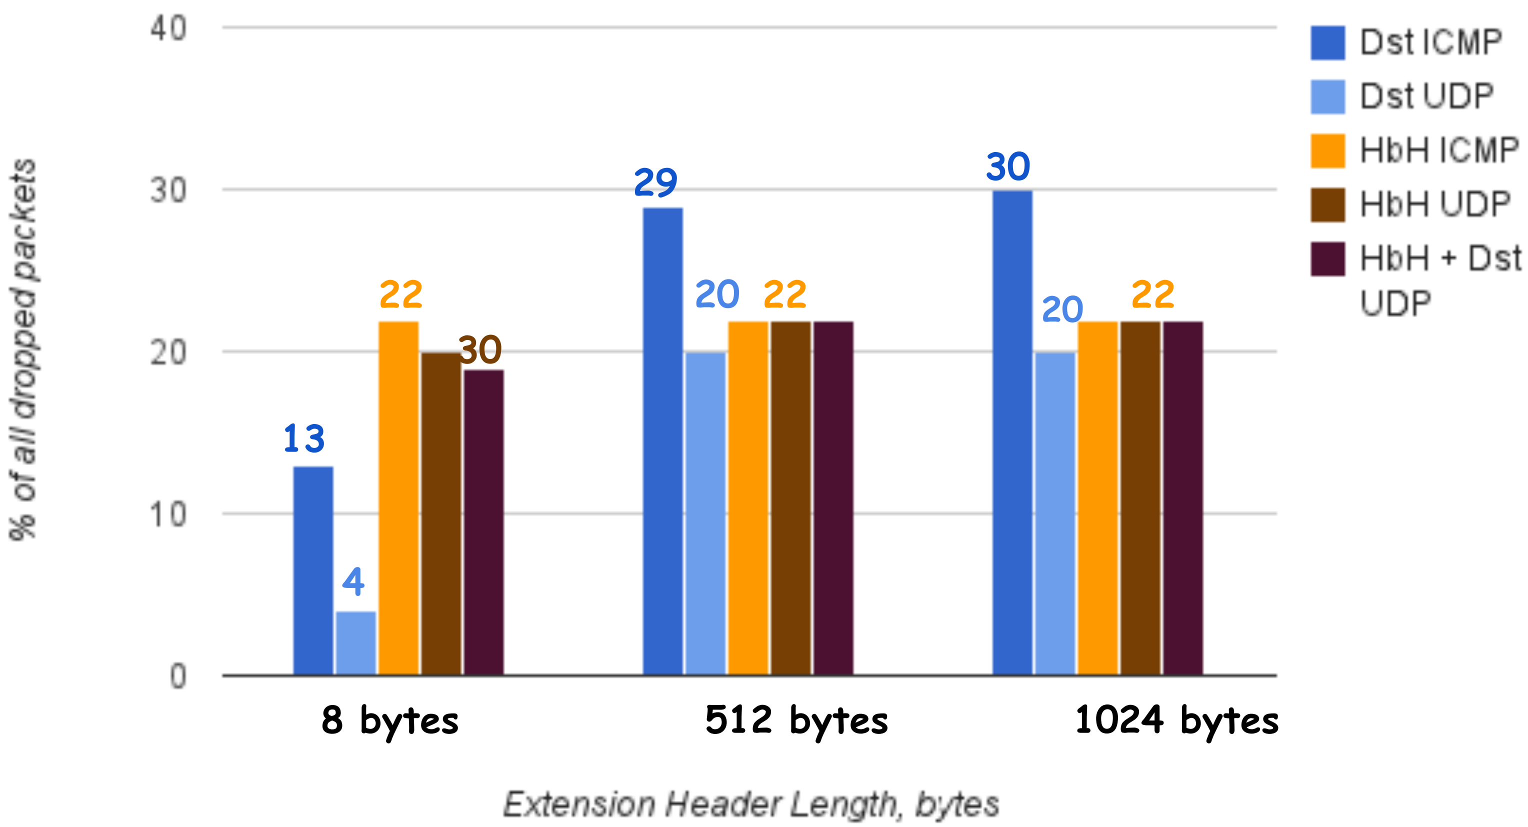

WIPv6LD dataset: ESP Drop rate

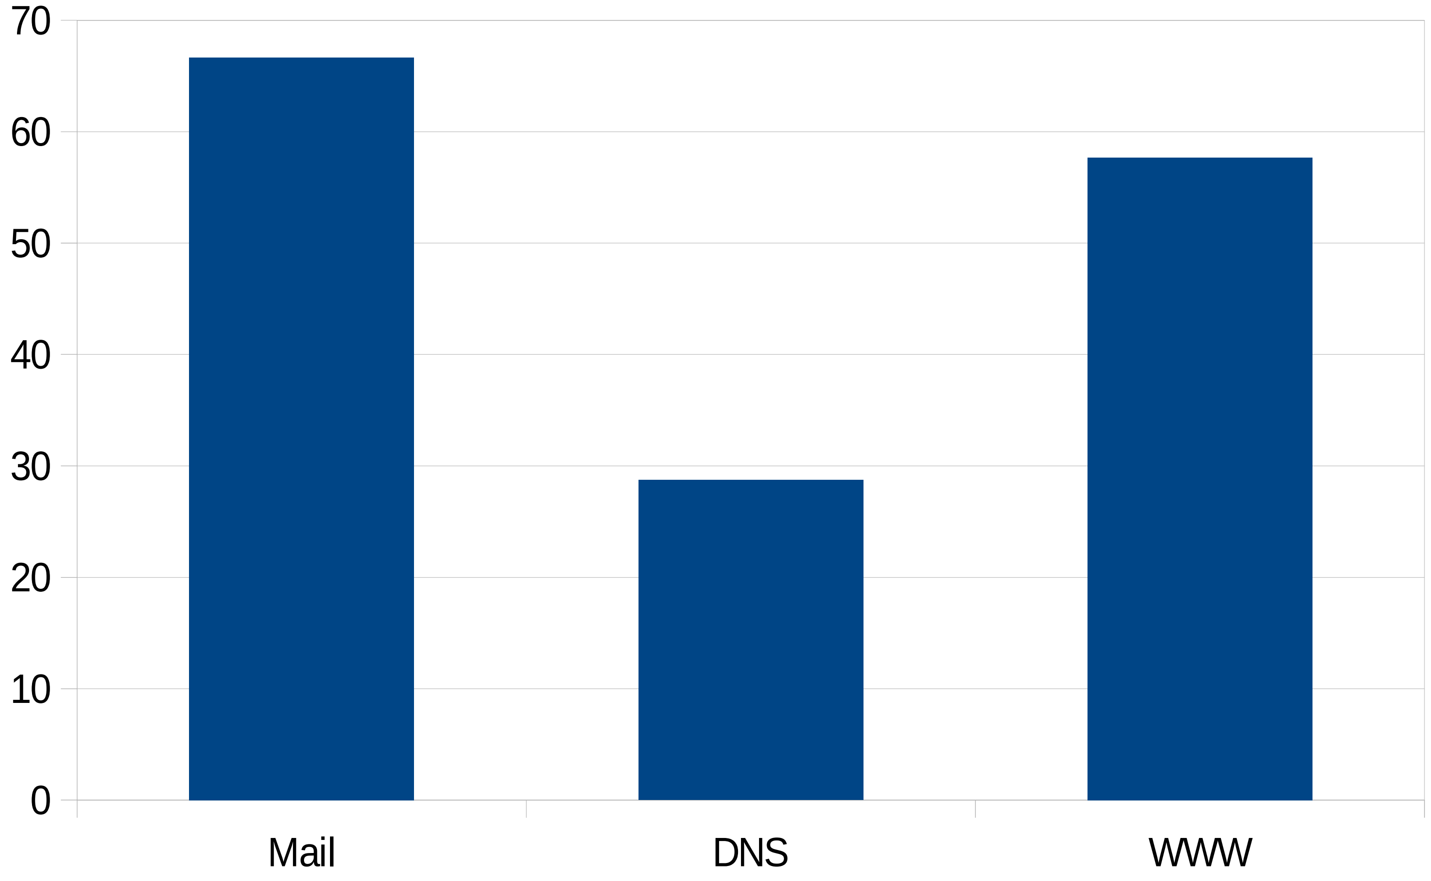

~~Speculations~~ Conclusions

- Packets with EHs ARE DROPPED ;(
- Short EHs have lower drop rate
  - most chips could not look deeper than first 64-128-256 bytes?
- For long EHs the next protocol does not matter
  - ACLs could not match it
- UDP packets with 8-bytes DO have the best chances to reach the destination
  - 80% success
  - ~50% of filtering - at the destination

Attacks with IPv6 Extension Headers

Old/obvious/boring stuff

- e.g. RA-Guard evasion

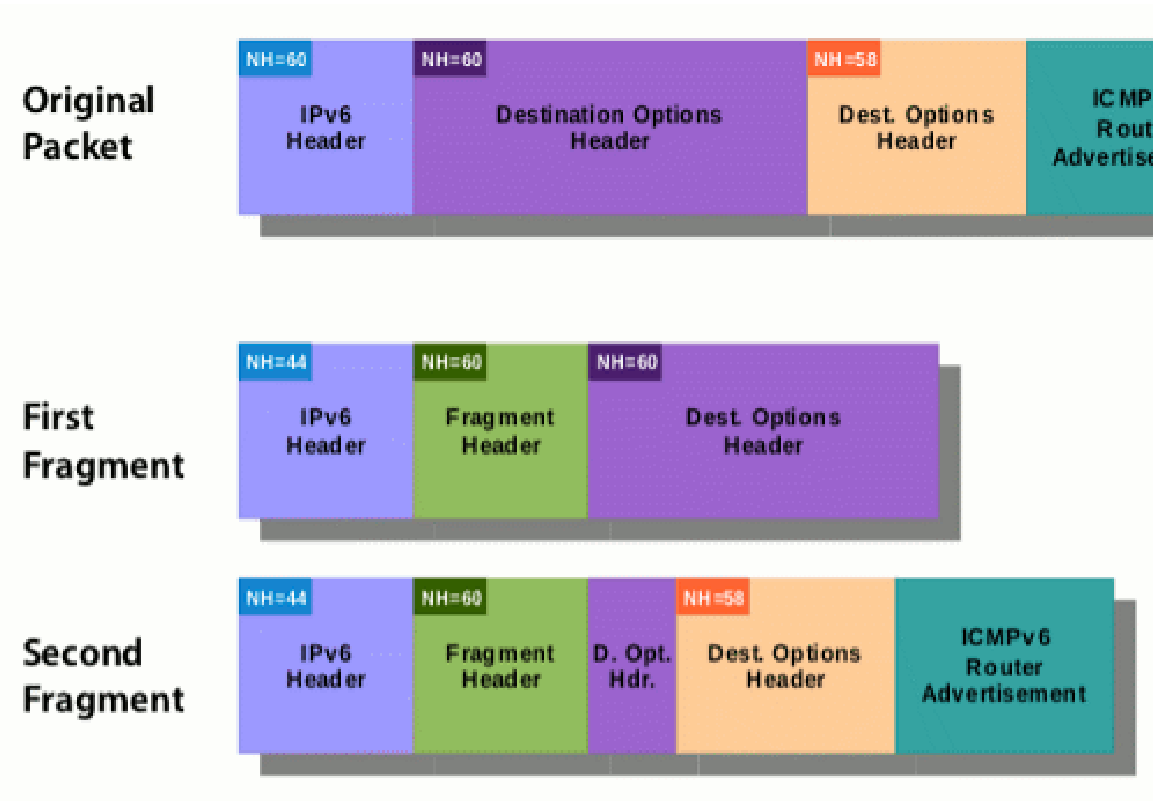

Original Packet

NH=60 IPv6 Header | NH=60 Destination Options Header | NH=58 Dest. Options Header | ICMPv6 Router Advertisement

First Fragment

NH=44 IPv6 Header | NH=60 Fragment Header | NH=60 Dest. Options Header

Can only tell there's Dest. Opt. Hdr inside!

Second Fragment

NH=44 IPv6 Header | NH=60 Fragment Header | D. Opt. Hdr. | NH=58 Dest. Options Header | ICMPv6 Router Advertisement

Can only tell there's Dest. Opt. Hdr inside!

More interesting stuff

- If IPv6 frags are widely dropped...What if we triggered their generation?
  - Send an ICMPv6 PTB with an MTU<1280
  - The node will then generate IPv6 atomic fragments
  - Packets will get dropped

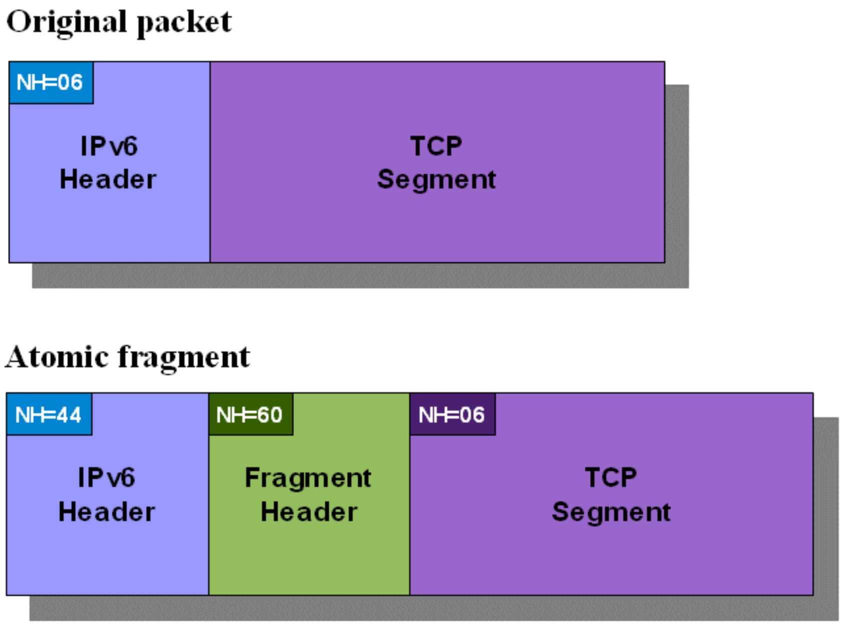

Attack Scenario #1

- Client communicates with a server

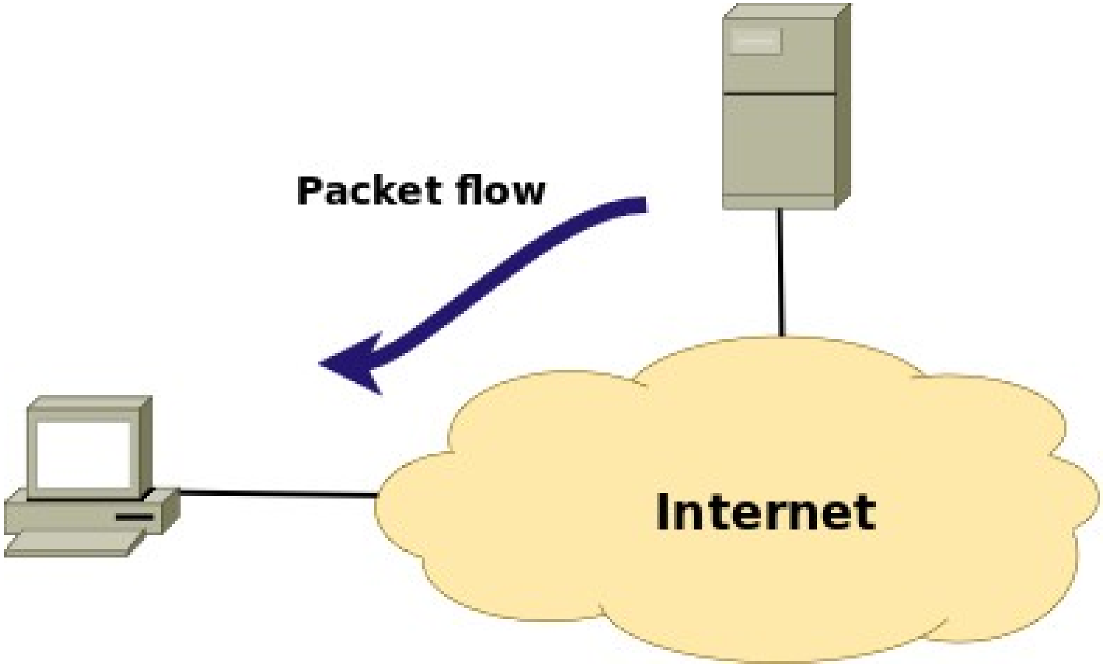

Attack Scenario #1 (II)

- Attacking client-server communications

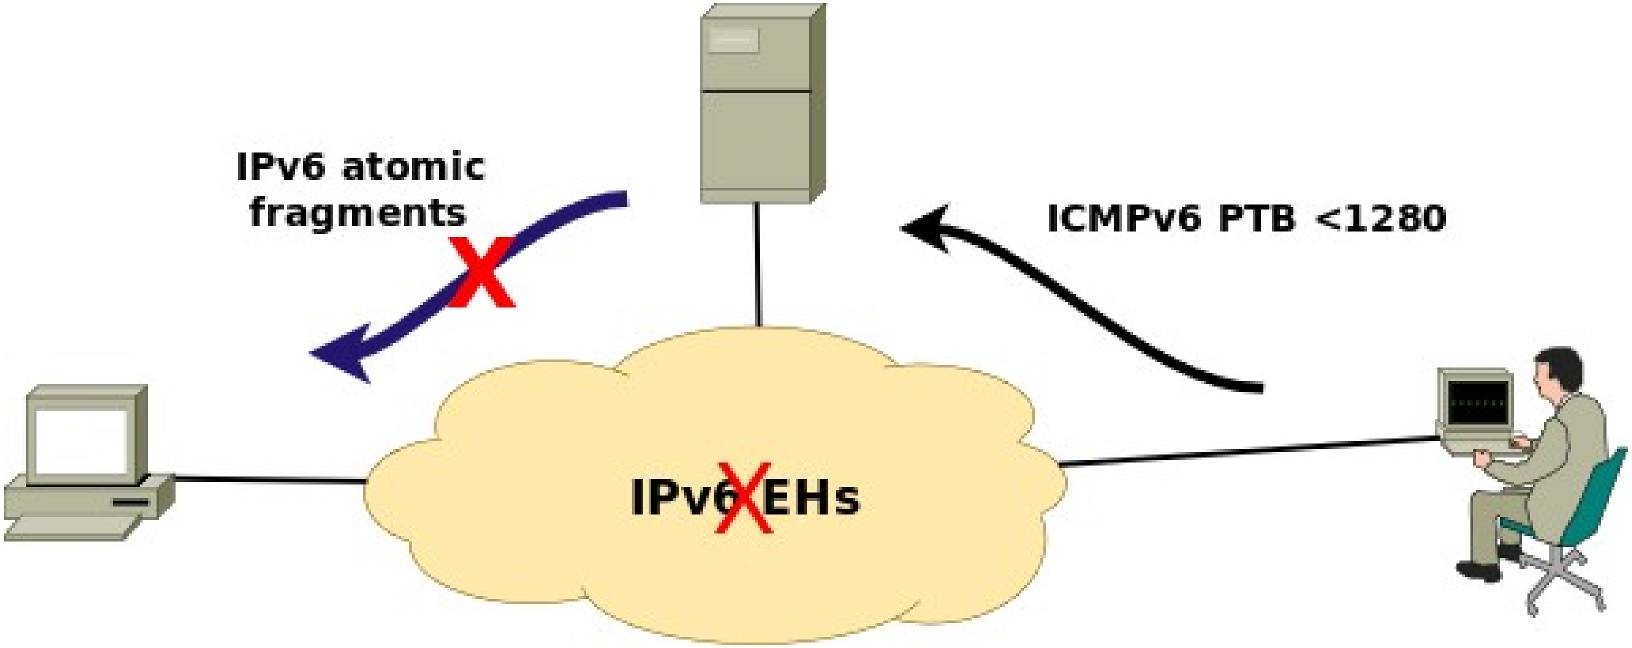

Attack Scenario #1 (II)

- Simple way to reproduce it:
  - Attack and client machine is the same one
  - So we attack our own “connections”
- Attack:
  - Test IPv6 connetivity:

    ```
    telnet 2001:4f8:1:10:0:1991:8:25 80
    ```

  - Send an ICMPv6 PTB < 1280 to trigger atomic fragments

    ```
    sudo icmp6  --icmp6-packet-too-big -d
    2001:4f8:1:10:0:1991:8:25 --peer-addr
    2001:5c0:1000:a::a37 --mtu 1000 -o 80 -v
    ```

  - Test IPv6 connectivity again:

    ```
    telnet 2001:4f8:1:10:0:1991:8:25 80
    ```

Attack scenario #2: BGP?

- Say:
  - We have two BGP peers
  - They drop IPv6 fragments “for security reasons”
  - But they do process ICMPv6 PTBs
- Attack:
  - Fire an ICMPv6 PTB <1280 (probably one in each direction)
- Outcome:
  - Packets get dropped (despite TCP MD5, IPsec, etc.)
  - Denial of Service

Mitigating these issues

- **draft-gont-6man-deprecate-atomfrag-generation**
  - "Do not send IPv6 atomic fragments in response to ICMPv6 PTB < 1280"
  - Update SIIT (IPv6/IPv4 translation) such that it does not rely on them
- **draft-gont-opsec-ipv6-eh-filtering**
  - Advice on filtering IPv6 packets that contain IPv6 Extension Headers

?

Questions?