Deprecating the Generation of IPv6 Atomic Fragments (draft-ietf-6man-deprecate-atomfrag-generation-00)

Fernando Gont
Will Liu
Tore Anderson

IETF 91
Honolulu, Hawaii, U.S.A. November 9-14, 2014

IPv6 atomic fragments

- RFC 2460 notes that when an ICMPv6 PTB < 1280 is received:
  - the node need not reduce the size of its packets
  - the node must generate atomic fragments so that the IPv6-to-IPv4 translating router can obtain a suitable Identification value to use in resulting IPv4 fragments

IPv6 atomic fragments & SIIT

- Atomic fragments could be of help when:
  - An IPv6 node communicates with an IPv4 node (through SIIT)
  - The IPv4 node is located behind an IPv4 link with an MTU < 1260
  - ECMP routing with more than one translator are employed for e.g., redundancy purposes

IPv6 atomic fragments & SIIT (II)

- Relying on IPv6 atomic fragments implies reliance on a number of factors:
  - ICMPv6 PTB being generated and delivered
    - PLPMTUD moved away from that for a reason
  - Generation of IPv6 atomic fragments in response to ICMPv6 PTB messages
    - Several OS/versions fail to do this
  - Fragments getting through to the intended destination
    - This results in additional fragmented traffic, which might get dropped

IPv6 atomic fragments & SIIT (III)

- Whether selecting the Frag ID at the source is an improvement is questionable:
  - The high-order 16-bits of the IPv6 Frag ID are stripped off
  - Relying on IPv4 reassembly at “high” data rates and IPv4 Frag ID's uniqueness has already been analyzed in RFC4963 and RFC6864
- Additionally, IPv6 atomic fragments offer an attack vector, even in non-SIIT scenarios (see next slide)

Attacks with atomic fragments

- Client communicates with a server

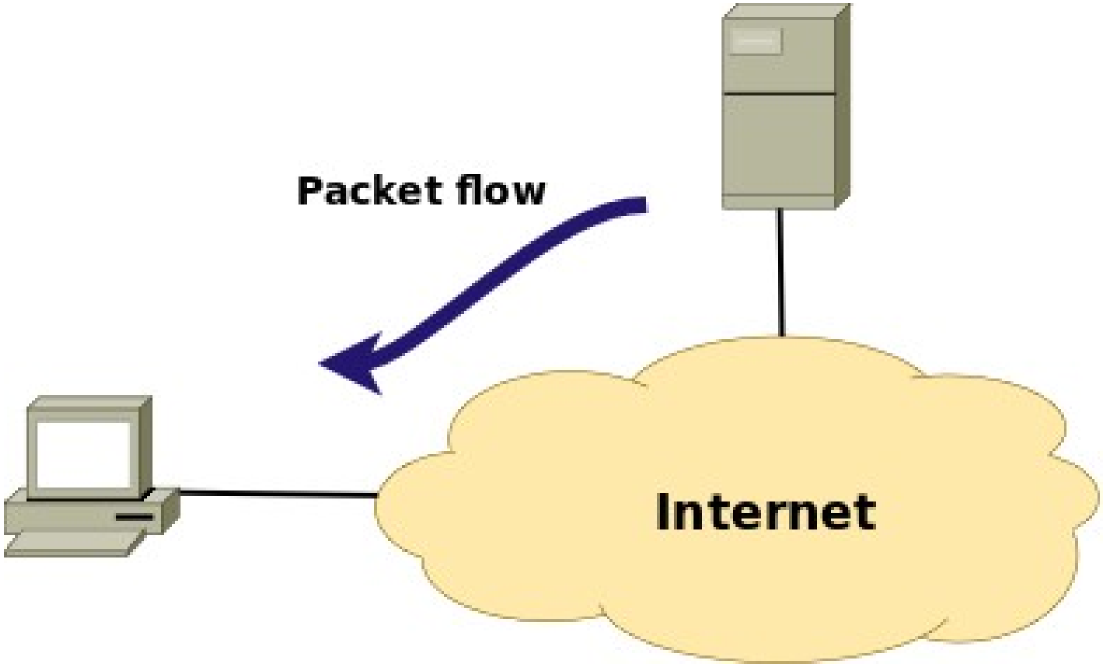

Attacks with atomic fragments (II)

- Attacker triggers atomic fragments
- Network filters IPv6 fragments

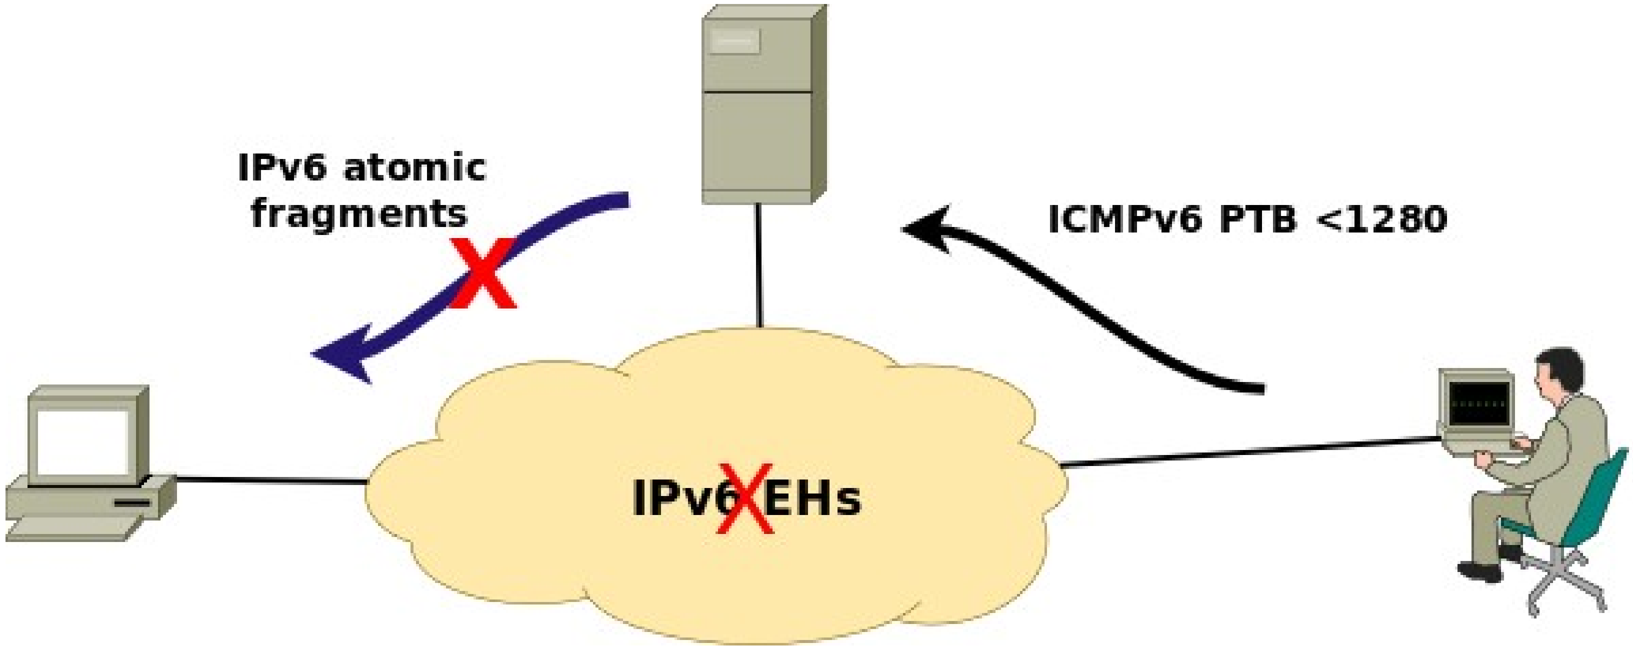

draft-ietf-6man-deprecate-atomfrag-generation

- Updates RFC 2460 such that:
  - atomic fragments are NOT generated in response to ICMPv6 PTB<1280
- Updates RFC 6145 such that:
  - The IPv6 FH is *not* employed to signal whether an IPv4 sender allows fragmentation
  - The MTU of an ICMPv4 -> ICMPv6 translated PTB message is never set to a value < 1280 bytes
  - When a resulting IPv6->IPv4 is <= 1260, the IPv4's DF bit is cleared

Comments?