Too Big or Too Small? The PTB-PTS ICMP-based Attack against IPsec Gateways

Ludovic Jacquin

Vincent Roca

Jean-Louis Roch

Globecom'14, Austin, December 2014

Packet Too Big (PTB) or Packet Too Small (PTS)?

The underlying idea…

About packet sizes and tunnel

- two gateways establish an IPsec tunnel to connect two remote LANs (or sites)

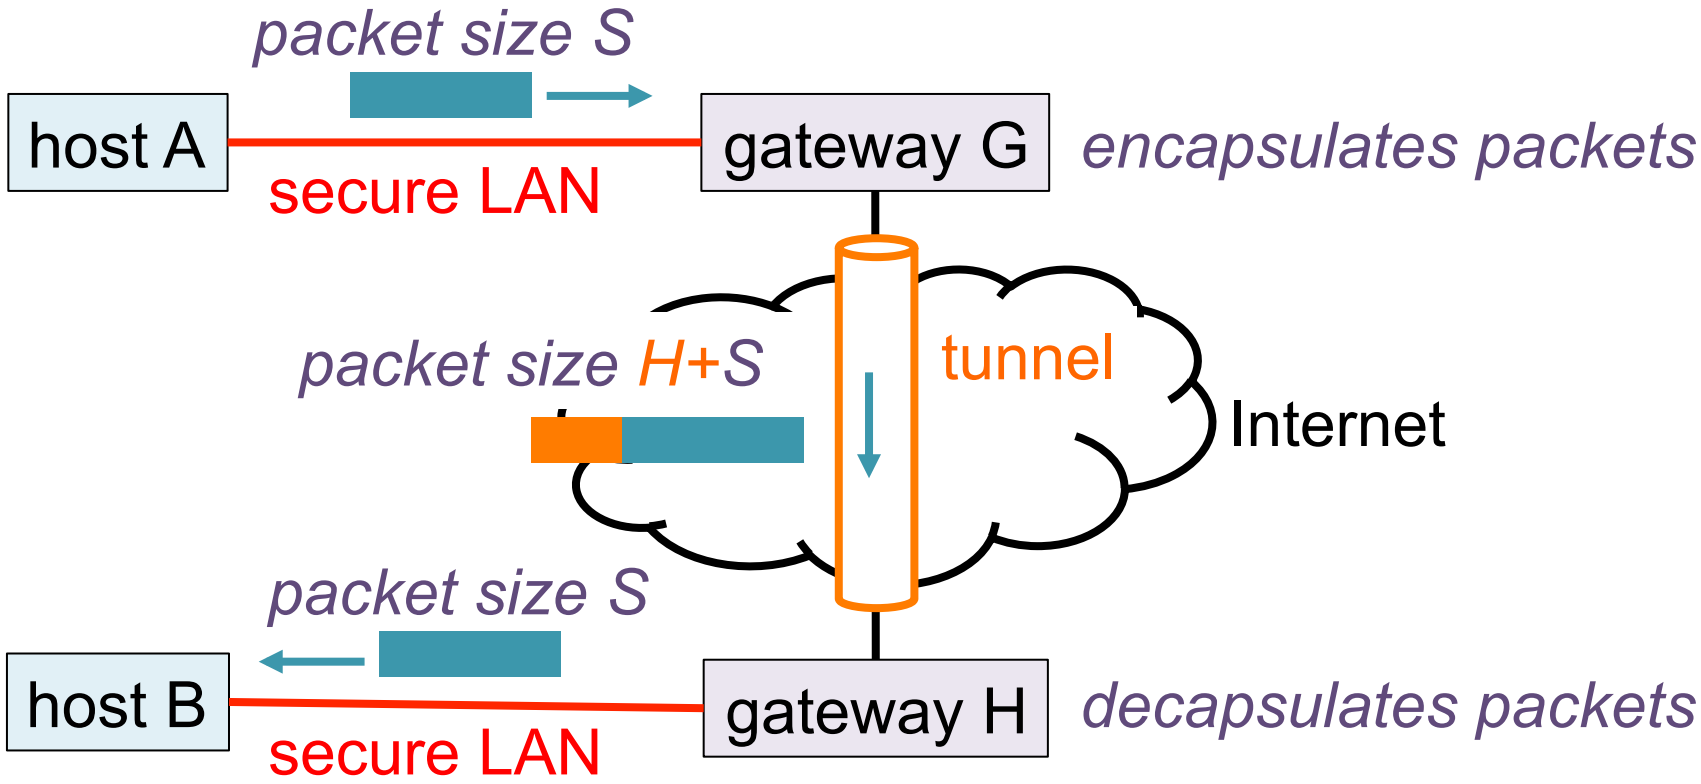

About packet sizes and tunnel… (cont')

- each link has a Maximum Transmission Unit (MTU)
  - **maximum allowed frame size on that link**
  - **e.g. 1500 bytes for Ethernet (i.e., 1460 b. or less at TCP level)**
- Path MTU (PMTU) is the min. MTU along the path
- a packet larger than a link's MTU is either
  - **dropped and an error ICMP "Packet Too Big" (PTB) message containing the MTU is returned to sender, or**
  - **fragmented if feasible (iff. IPv4 with DF bit clear)**
- each link MUST guaranty a minimum MTU
  - **IPv4 576 bytes**
  - **IPv6 1280 bytes**
  - **essentially here for performance reasons**

The issue

- what happens if G's outgoing link is already at MTU 576 bytes (IPv4)?
  - then we need H+S ≤ 576, which implies that S < 576…

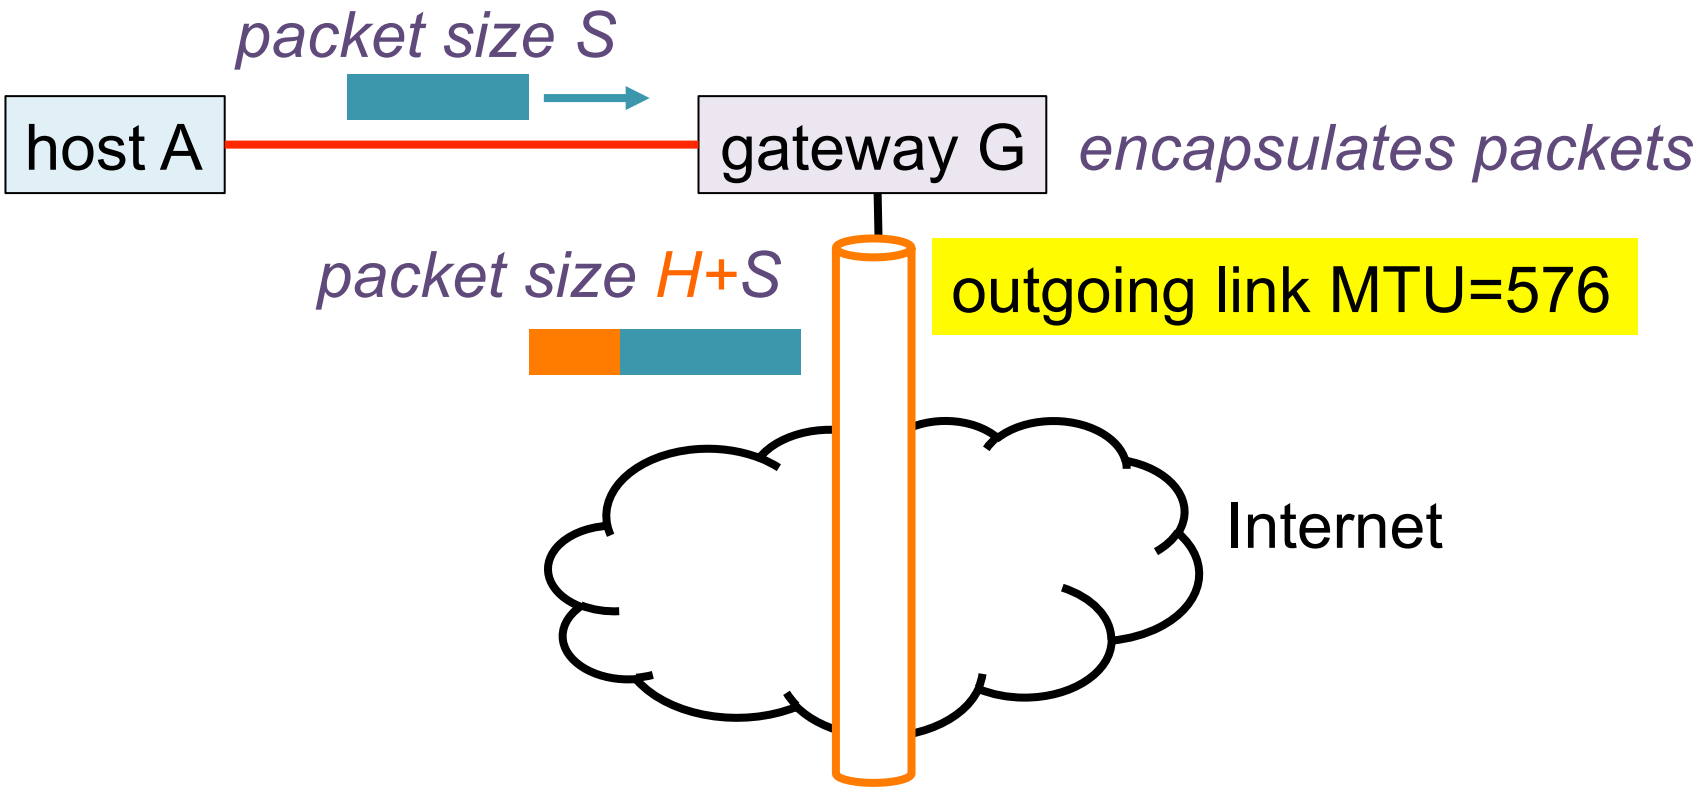

The issue… (cont')

- we observed, through experiments, that A and G don't understand each other

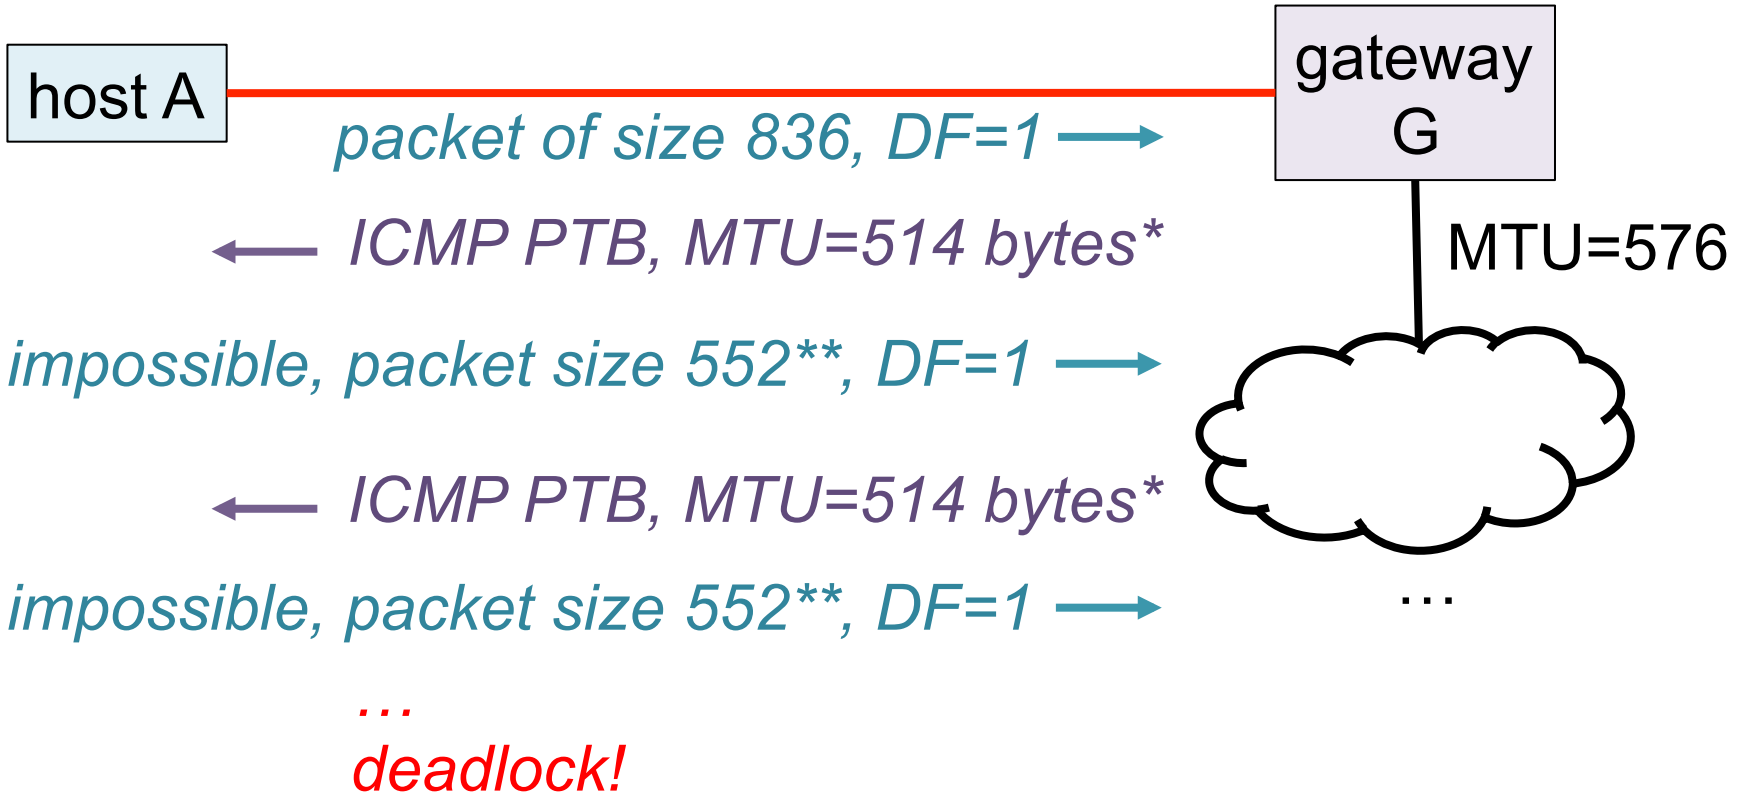

* 514 bytes because of IPsec ESP header

** 552 is minimum PMTU value on Debian/Linux

The issue… (cont')

- the reality is slightly more complex…

  - does A use:
    - **PMTUd (Path MTU discovery) (default)**
      - based on probing with DF=1, listening to ICMP PTB
    - **PLPMTUd (Packetization Layer PMTUd)**
      - TCP-level (or similar) probing mechanism, taking advantage of TCP ACK. ICMP PTB messages are totally ignored
  - is it a TCP or UDP flow?
    - no delivery guaranty with UDP!
  - is it IPv4 or IPv6?
    - IP fragmentation prohibited from IPv6

Details of our IPsec/ESP exploit

Description of exploit

- IPsec configuration based on-the-shelf components
  - **stable Debian "Squeeze" distribution**
  - **end-host, gateway and IPsec default configuration**

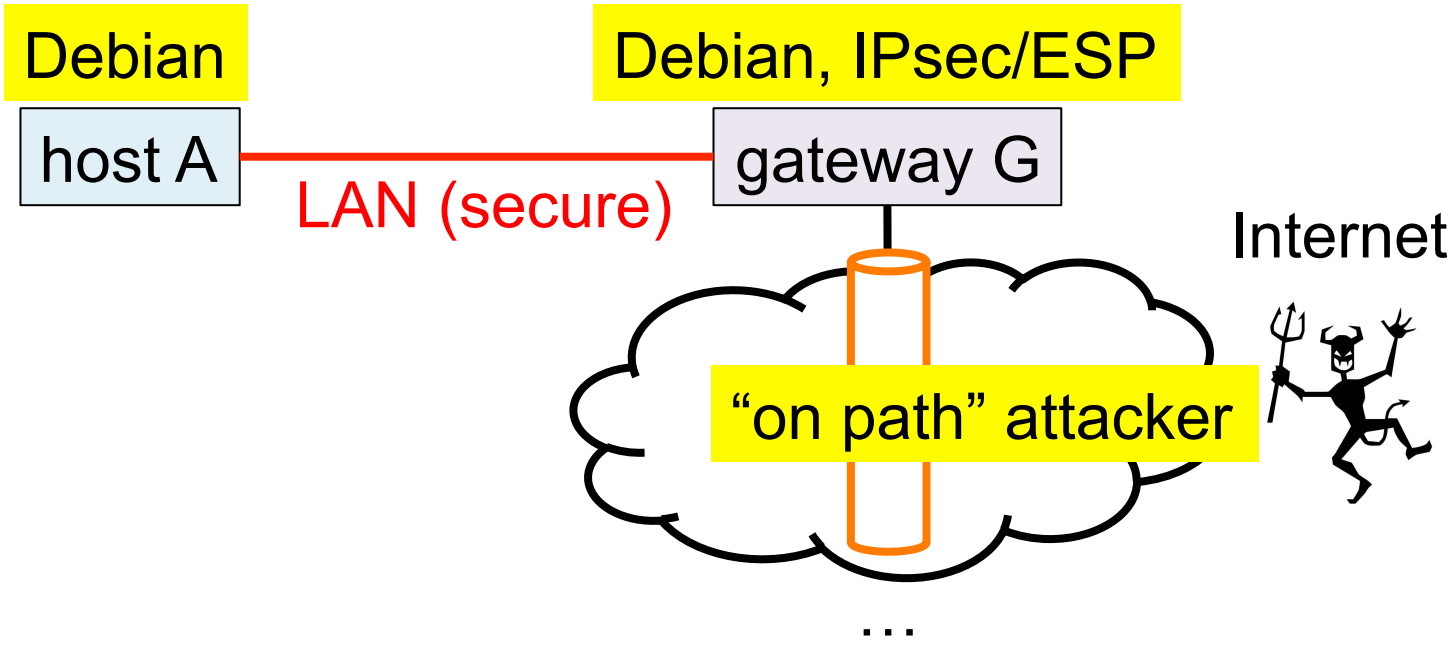

Description of exploit… (cont')

- launching the attack to gateway G
  - the attacker needs to **be on the IPsec tunnel path**
    - **eavesdrops a tunneled packet, then**
    - **forges an ICMP PTB ("Pkt Too Big") message that contains a copy of the eavesdropped packet**
      - needed to bypass IPsec security WRT. ICMP error messages
  - the attacker can be a compromised router…
  - … or a simple host attached to a **non-encrypted WiFi**
    - **if a user uses an IPsec VPN to his/her home network, and is attached to this non-encrypted WiFi, then we can attack the remote IPsec gateway**
  - a **single** "well formed" ICMP PTB packet is sufficient to launch the attack!

TCP/IPv4, PMTUd configuration

- on host A, TCP fetches the local MTU updated by the PMTUd, but does not go below 552 bytes

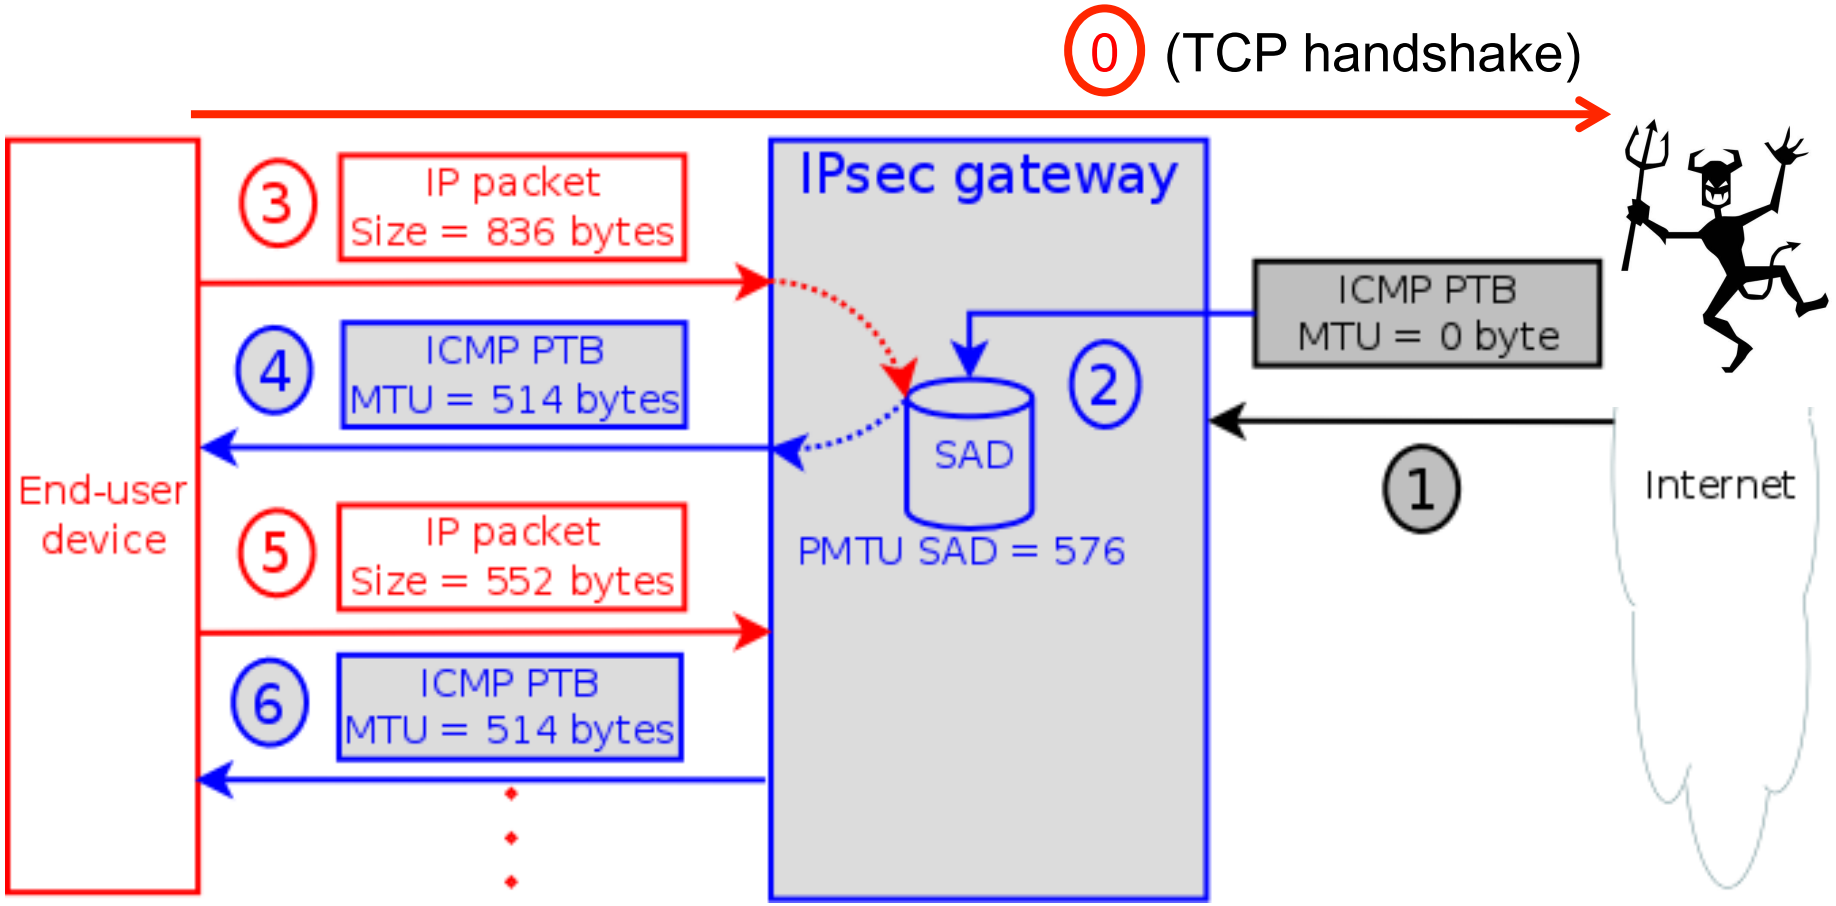

TCP/IPv4, PLPMTUd configuration

- on host A, TCP ignores local host MTU configuration and relies on TCP-level PLPMTUd

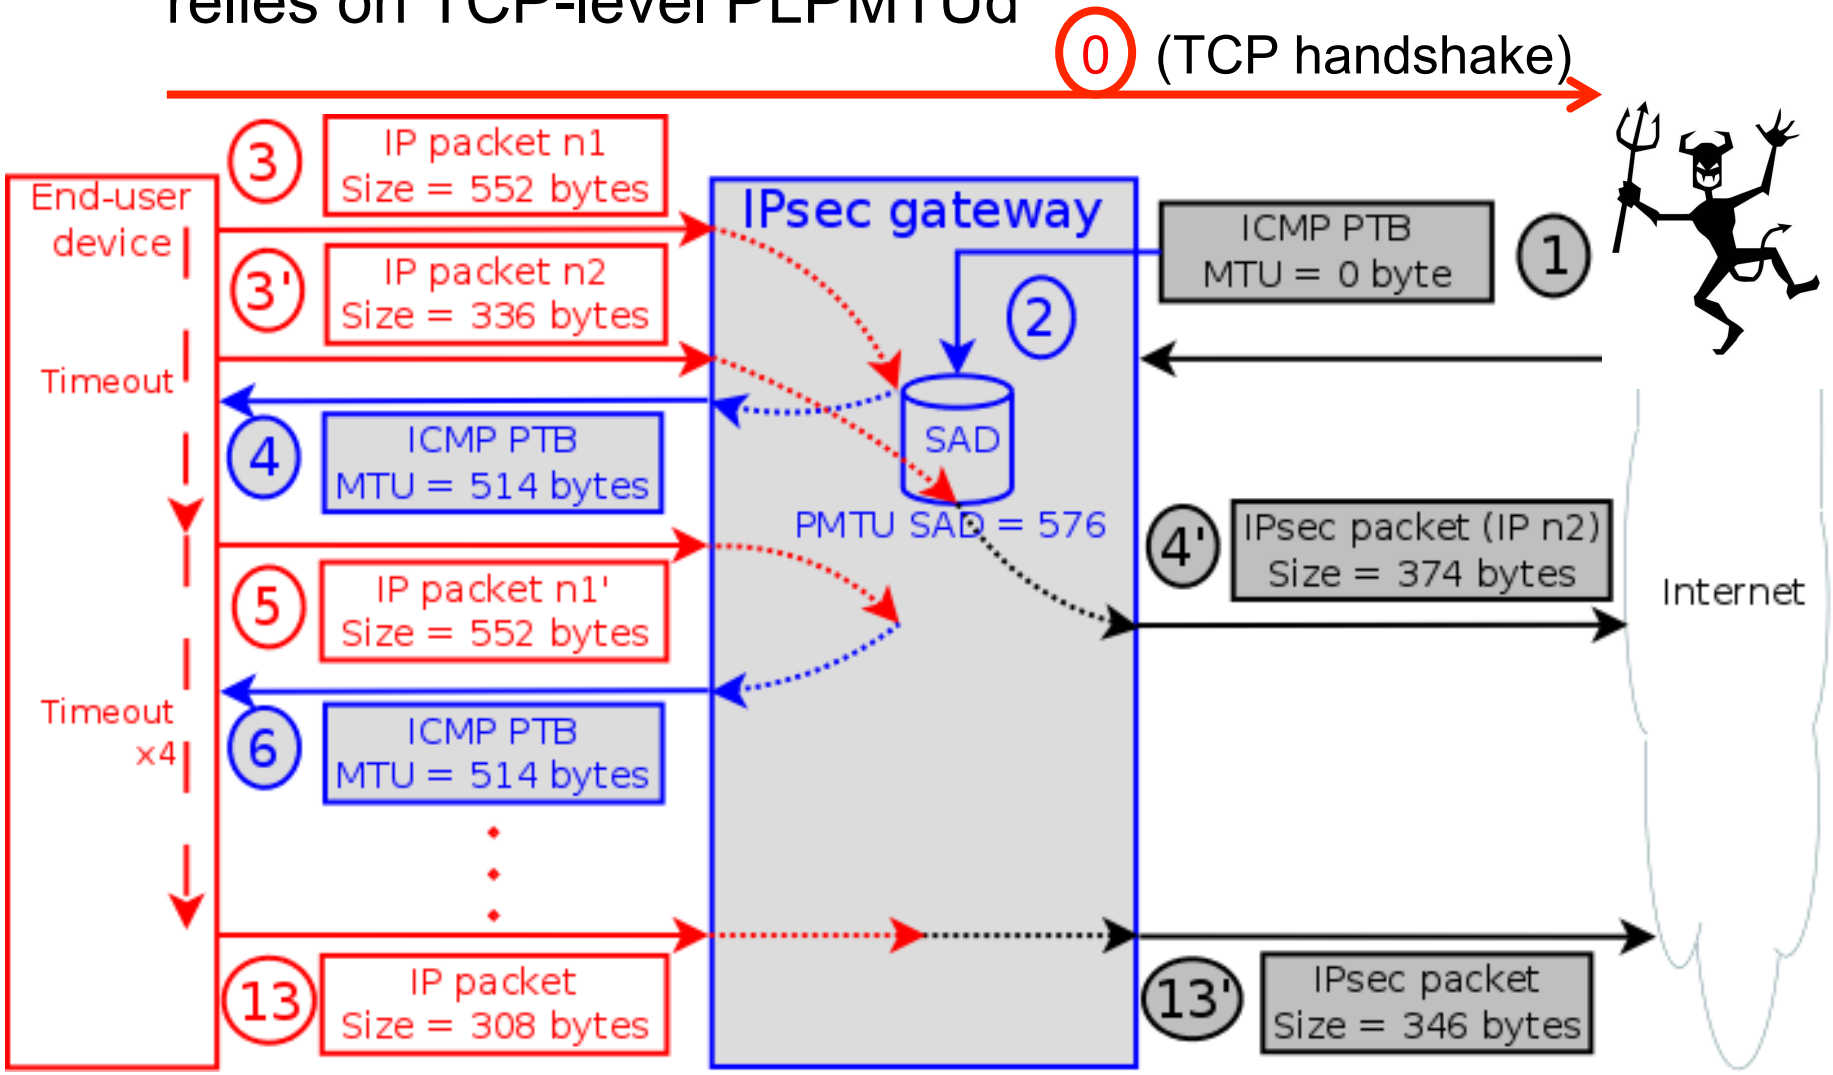

Results

TCP, IPv4, PMTUd	**DoS:** no connection possible any more
TCP, IPv4, PLPMTUd	**major performance impacts:** 6.5s initial freeze, then tiny packets
UDP, IPv4, PMTUd	**major performance impacts:** tiny packets
*, IPv6, *	not tested

- yes, it works pretty well, with impacts that depend on the exact configuration

Conclusions…

To conclude

- a highly effective attack
  - a **single** packet is sufficient to launch it
- the problem may be more serious than just an IPsec DoS
  - **ambiguity** in the way minimum MTU should be handled in presence of a tunnel, no matter the tunnel type
    - **To be confirmed**

To conclude… (cont')

- which solution to the problem?
  - gateway G **should not accept ICMP** error feedbacks?
    - **don't agree, ICMP is useful per se and so is PMTUd in highly dynamic networks to find the right PTMU in a Layer-4 independent way**
  - gateway G **should not accept** MTU=576 as it knows it's incompatible with tunneling?
    - **YES, but what about MTU=676, just a little bit larger?**
    - **well, it will be accepted, and still negatively impact performance, even if a less severe way (no DoS)…**
  - gateway G **should be able to explain host A** that using a lower value than 576 is valid in this case?
    - **YES, but it remains tricky… What if there's a 2nd tunnel?**

To conclude… (cont')

- which solution to the problem… (cont')
  - gateway G **should always be able to fragment**?
    - **even if DF=1? Even with IPv6? MAY BE… but it's tricky!**
    - **BTW, there's a Cisco IOS 12.2(11)T note explaining DF should be ignored!**
  - ICMP PTB error messages coming from Internet should be confirmed with a separate mechanism
    - **could be a probing scheme similar to what PLPMTU does**
    - **of course a powerful attacker on the path could identify these probes and drop them…**
    - **but an active attack that modifies the flow is easier to spot!**
  - that's for future work…