Generation of ICMPv6 Echo Replies for Teredo Clients

draft-denis-icmpv6-generation-for-teredo-01

Teemu Savolainen,
Rémi Denis-Courmont

Behave WG at IETF #77

March 2010

NOKIA

Teredo and ICMPv6

- Teredo, as per RFC4380, uses return routing and ICMPv6 to discover the closest Teredo relay corresponding to any given peer
- Unanswered ICMPv6 Echo Requests make connection creation fail as Teredo client assumes peer is unreachable
- ICMPv6 Echo Request/Reply is assumed to work through Internet, if a peer is reachable

Working NAT64 scenario

1. A Teredo client sends an ICMPv6 echo request via Teredo Server
2. The ICMPv6 echo request is translated to an ICMPv4 echo request at a NAT64
3. An ICMPv4 echo reply is translated to an ICMPv6 echo reply at the NAT64
4. The ICMPv6 echo reply is received successfully at the Teredo client via a Teredo relay, of which IPv4 address and used UDP port the Teredo client thereby discovers and can use for IPv6 communciation with the peer

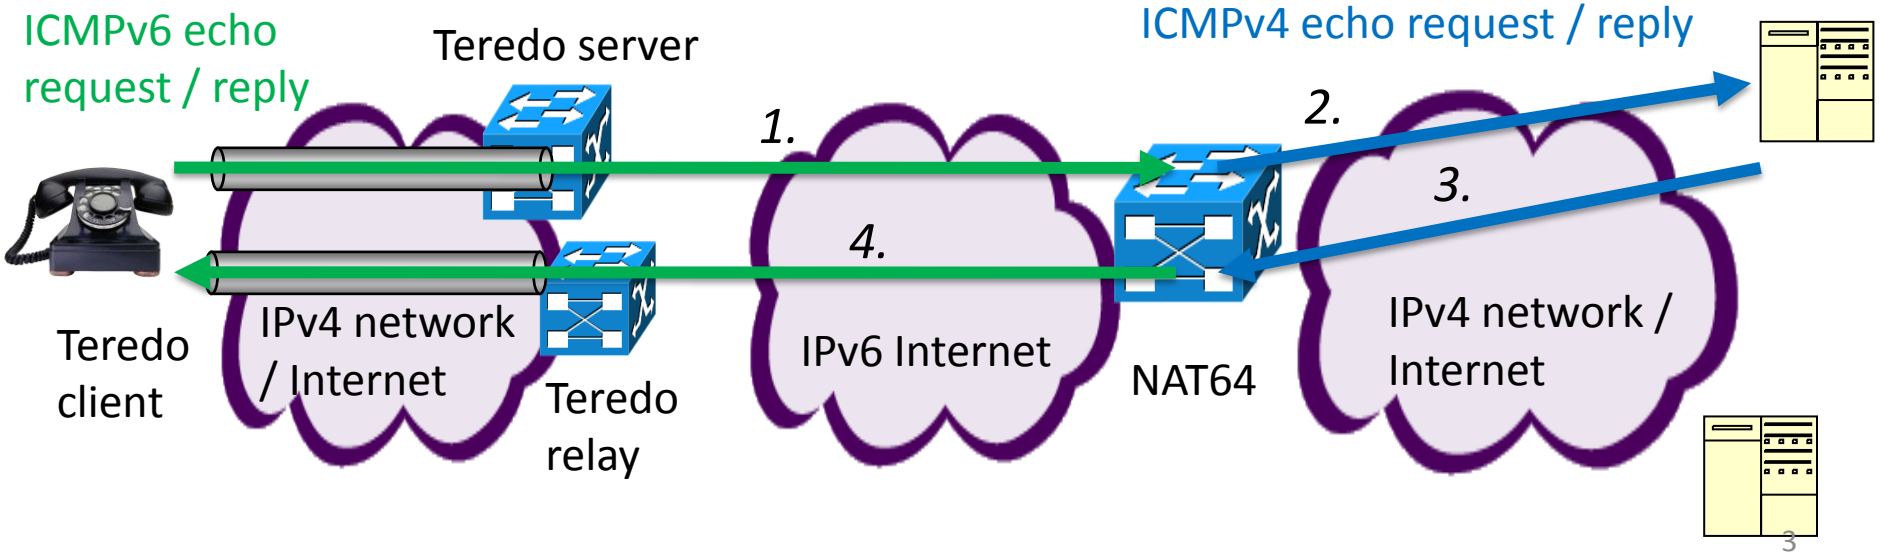

Broken NAT64 scenario

1. A Teredo client sends an ICMPv6 echo request via Teredo Server
2. The ICMPv6 echo request is translated to an ICMPv4 echo request at a NAT64
3. An ICMPv4 echo reply never arrives at the NAT64 for **some** reason
4. The Teredo Client never receives ICMPv6 echo reply, and thus cannot continue its connection creation process, as it does not discover Teredo relay, and as by design the Teredo server cannot be used to carry IPv6 traffic

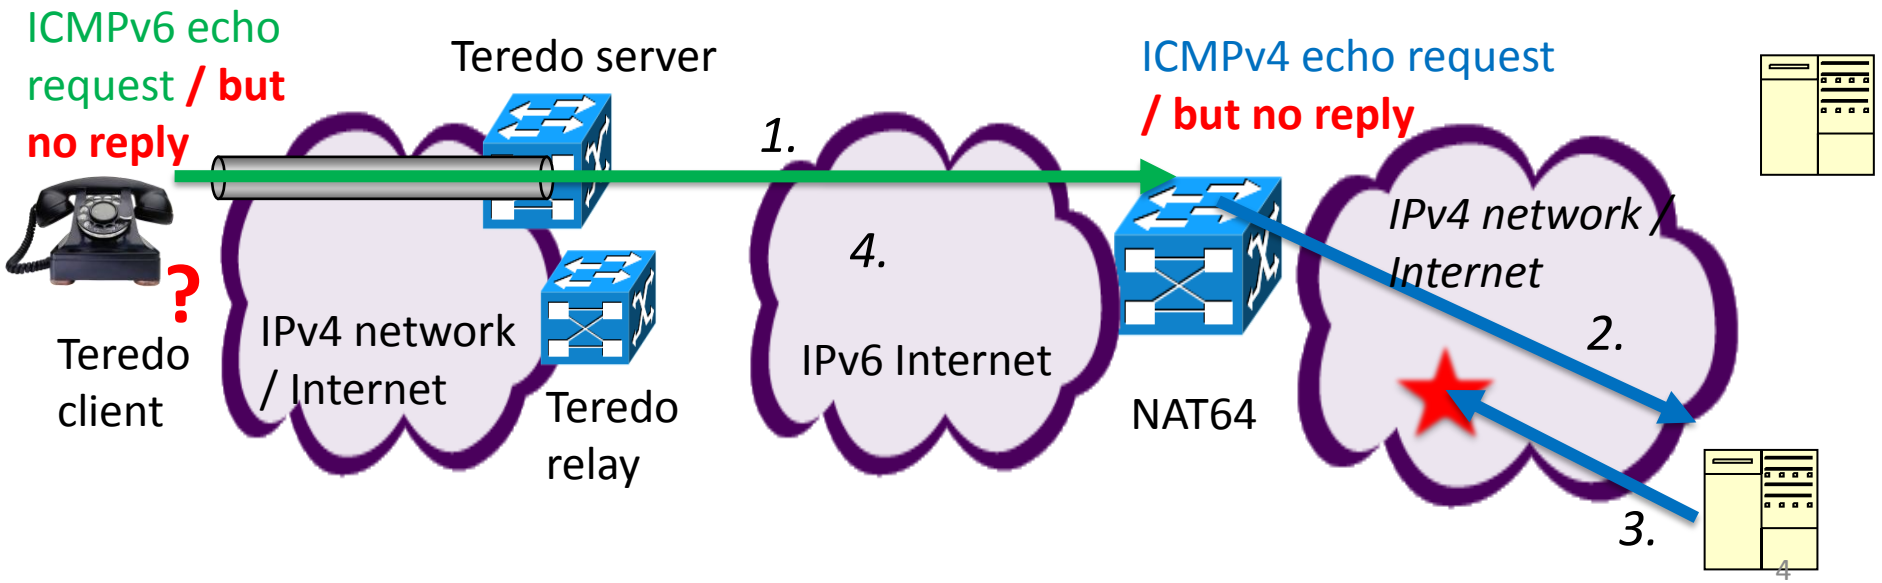

Likelyhood of the scenario

- According to Google's statistics, revealed at March 2009 workshop, **2%** of their IPv6 users were using Teredo
  - Interestingly, **1.4%** of Google's IPv6 users were shown to use Teredo in Google's presentation at October 2008 (increasing trend perhaps?)
- The problem was accidentally found in an artificial network setup, not in a real network – except real IPv4 sites did not respond to ping
- If this Teredo population would access these IPv4 sites through NAT64, they could hit the problem...
- No problem reports heard from live networks this far

=> A problem, but not very likely one

- Statistics here: https://sites.google.com/site/ipv6implementors/conference2009/agenda/10_Lees_Google_IPv6_User_Measurement.pdf?attredirects=0 and http://www.ripe.net/ripe/meetings/ripe-57/presentations/Colitti-Global_IPv6_statistics_-_Measuring_the_current_state_of_IPv6_for_ordinary_users_.7gzD.pdf

Way forward?

- Is Behave WG interested to work on this problem (e.g. when occurring in NAT64 scenario)?
- Or should the work and the draft be shelved until the problem starts to appear in real life *(if ever)*?