The Trouble with Transports

Moving UDP to layer 3.5

IAB:
Brian Trammell
Joe Hildebrand

Things we need to admit

- TCP isn't enough
  - e.g. congestion, forced retransmit, rigid 3-way handshake
  - HTTPS-over-TCP is even more limiting
  - Minion is cool but requires kernel mods
  - Some L2 protocols do not interact well (LTE)
- SCTP won't save us
  - 15 years of proof
  - Lots of middleboxes, kernels, APIs to fix
- NATs will continue to exist
- New protocols on UDP: WebRTC(SCTP), QUIC, RTMFP, MOSH, etc.
- Conclusion: nothing new at layer 4

Goals

Endpoints

- Expose minimum information to get midpoints to allow traffic
- Transport/application innovation
- Maintain/increase scalability
- Explosion of local addresses
- Hints from/to network path
- Detect broken paths?

Midpoints

- Characterization on fast path
- More explicit policy than possible with "plain" UDP
- Less need for rapid evolution with new transports
- Add more explicit value
- Limit traffic evasion through 443/tcp by being **reasonable**

Design Constraints

- Doable in user space only
- Existing socket API to kernel
- No root privileges
  - No raw sockets
  - No "privileged" ports (nice try, 20th century)
  - No ICMP
  - No access to DHCP or RA info
- Multiple apps per endpoint
- No worse than TCP
  - Privacy
  - Channel cohesion
  - DoS resiliance
  - Hope to do better
- Existing UDP middleboxes must work without modification
- Stable over time
- Identification might be probabilistic

Hopeful assumption: middleboxes that block UDP are the ones that get updated frequently (i.e. corporate firewalls)

General approach:
Protocol inside UDP: "SPUD" for now

- One bidirectional relationship per endpoint pair
- Multiple relationships per port
- Explicit signaling from/to path/endpoints
- Stay async wherever possible
- Leave everything possible to new-transport layer
  - Retransmit
  - Congestion control
  - Separate security properties for network and endpoint

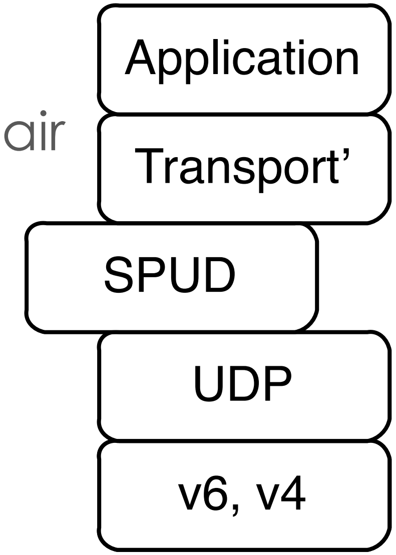

Interactions with other IETF work

- TAPS: determine use cases for transports *inside*
- AECON: possible protocol
- APONF: unknown
- RTCWeb: potential firewall traversal for future
- RMCAT: input to congestion control for transports
- DTLS: potential requirements for v3 as implementation…
- Likely many others