RTP Media Neutrality

Harald Alvestrand
Speaking for myself

Background

- draft-alvestrand-rtp-sess-neutral
- Ca 1994: MIME types and RTP media types merge their namespaces (RFC 1889 was 1996)
  - Conversation with Steve Casner on a Prague subway
  - RTP RFC 1889: 1996
  - SDP RFC 2327: 1998
- Almost immediately (within 10 years), tensions became obvious
  - MIME: video/vp8; parameter=value
  - RTP: m=video / a=rtpmap:99 vp8/90000 / a=fmtp:99 .....

Flows vs Transports

(not using the word "multiplexing" - it confuses)

- A media flow can be interpreted alone* by a decoder
- A transport pair can be told apart by intermediate nodes

Historical changes, 1996 to today

- Cost of transport pairs go up
- Range of flows increases
- Cost of bandwidth plummets
- RSVP and multicast do not deploy

* modulo repair flows, FEC flows and so on

Help from the network

- Prioritization of flows
  - Requires detecting them
    - 5-tuples, DSCP markings, Deep Packet Inspection
  - "Do this one first" or "Reserve capacity for this one"
    - Has no effect when capacity is plentiful
  - (Censorship is one version of "prioritization")
- And...?

The RFC 3550 leap of faith

- Network-based prioritization is important
  - Assumes resource shortfalls will be frequent
  - End-system prioritization doesn't need separate transports
- Media types neatly align with priorities
  - Video and audio needs to be separated
  - Even if priorities are equal, cost does not matter
- Having many transports is cheap
  - RFC 3550 even used destination addresses only. Today, a 5-tuple and a crypto key is usually assumed.

RFC 3550 Arguments - head on

Section 5.2 bullet 1-3 are irrelevant (SSRC reuse)

Bullet 4: "An RTP mixer would not be able to combine interleaved streams of incompatible media into one stream."

- Mixing is one application. "Blind" media mixing is appropriate for telephone conferences, and just about nothing else.
- In most cases, a mixer needs much metadata to do its job correctly. Sometimes, even interrelations between audio and video tracks are important for it to do its job.
- This argument's validity is strongly application dependent.

RFC 3550 section 5.2 bullet 5

"Carrying multiple media in one RTP session precludes: the use of different network paths or network resource allocations if appropriate; reception of a subset of the media if desired, for example just audio if video would exceed the available bandwidth; and receiver implementations that use separate processes for the different media, whereas using separate RTP sessions permits either single- or multiple-process implementations."

- Note the use of *if desired* here.
- Media types *do not align* with these desires.

Multiple RTP sessions are often desirable. Each of them may need audio, video, application data or all 3.

Corrective Action

- The user knows the tradeoffs
  - Let the user decide what transports to allocate
- RTP needs no change
  - RTP sessions and flows have no MIME type markers
- RTCP is probably OK
  - Feedback flows may get surprising timings for mixtures of different-sized flows
  - This problem occurs also within one media type
- SDP needs corrective action
  - draft-holmberg-mmusic-sdp-bundle-negotiation is a patch on the current syntax
  - This discussion is input to the design of SDP:TNNG (the one after SDP:TNG)

What next?