RTP Multiplexing

draft-rosenberg-rtcweb-rtpmux

Jonathan + {Rosenberg, Lennox}

Problem Definition

- Today RTP spec requires a separate RTP Session for different media types
- Each RTP session requires a separate port
- Using separate ports for audio and video means:
  - Double number of ports consumed in NAT – ports need to be preserved
  - Increase ICE setup time
  - Increase failure cases
  - Adds complexity

What we really want

- Single transport session (e.g., 5-tuple) for all p2p content exchanged between two rtcweb endpoints
  - Voice stream(s), video stream(s) over RTP
  - Data somehow muxed as well (approach TBD)
  - Single ICE negotiation
- The on-the-wire protocol still is RTP so it works with sniffers/firewalls/etc
- Framing and operation identical to usage with a single media stream
- Reuse existing RTP codebases

Solution: Use SSRC for demux

- RTP already allows for multiple SSRC in a session in several cases
  - Each participant is a unique SSRC sending a unique stream
  - A single participant can send multiple streams of the same media type (e.g., talking head video and screen-share video)
- We will naturally extend this by allowing a single participant to send multiple streams of different media types (e.g., audio and video)
- SSRC is still allocated through random selection and has no structure (NOT the structure proposed in draft-rosenberg-rtcweb-rtpmux)

Implications

- Since this is not backwards compatible, the signaling will need to negotiate its usage
  - This is an issue only when using webRTC to talk to other domains via SIP, or existing SIP devices
- The webrtc client must also support non-multiplexed voice and video for inter-domain and interop purposes, where each RTP session uses its own UDP flow
- If webrtc wishes to use FEC on the same transport, additional signaling needs to be defined to support it
- Combining media streams of different bandwidth on the same session will cause RTCP intervals to have odd timings; we will need to specify to act more sanely

What gets specified

- RTCWeb specification will:
  - Mandate implementation of muxing based on SSRC within the browser
  - Mandate support for using a different RTP session (and transport) for each media type
  - Mandate a set of parameters for RTCP transmissions that keep things sane (which probably means AVPF is required and used)
- A new SDP specification will be written which defines how to signal this multiplexing in SDP
  - Only needed by rtcweb if rtcweb includes SIP in the browser
- AVT will define a more complete solution which addresses all of the multiplexing cases well – not a dependency for rtcweb