University of Glasgow | School of Computing Science

WebRTC Media Transport and Use of RTP

draft-ietf-rtcweb-rtp-usage-10

Colin Perkins – University of Glasgow
Magnus Westerlund – Ericsson
Jörg Ott – Aalto University

Changes Since Last Meeting

- Changes in -08:
  - Rewrote Section 12 (“RTP Implementation Considerations”)
  - Removed most of the Appendix (“Supported RTP Topologies”), moving the remainder into Section 12
- Changes in -09:
  - Updated references
- Changes in -10:
  - Clarified that keying for RTP/SAVPF profile specified in security-arch draft
  - Clarified that an endpoint can have multiple RTCP CNAMEs if it sends streams synchronised to multiple clocks
  - Clarified that the RTP circuit breaker is a boundary condition, and that applications also need to implement congestion control
  - Clarified that RTP/AVPF + DTLS-SRTP keying is mandatory to implement

Open Issues

- Several open issues remain to discuss:
  - Signalling coding capability
  - Signalling RTP topologies
  - Simulcast
  - Forwarding media
  - Use of differentiated services
  - Mapping to W3C API
  - Correlating media streams

- Would like to resolve most of these this week
  - Some might be resolved by moving the discussion to separate drafts

Open Issue: Signalling Coding Capability

- Do endpoints need to signal limitations in their capability to encode or decode some number of simultaneous streams?
  - One possible proposal is in draft-westerlund-mmusic-max-ssrc-02
    - Defines media-level "a=max-send-ssrc:" and "a=max-recv-ssrc:" SDP attributes
    - Are media-level attributes sufficient when using the SDP bundle extensions?
  - Currently just require "support for use of multiple simultaneous SSRC values in a single RTP session" with no limit on the number of SSRCs or flows that can be encoded/decoded
  - Affects Section 4.1 and Section 12.1.1

Open Issue: Signalling RTP Topologies

- WebRTC endpoints use one or more RTP sessions in the context of a PeerConnection
    - Each RTP session can convey several RTP media streams, possibly from several capture devices, representing layered coding, or for FEC
    - Each RTP session can extend beyond the scope of single PeerConnection if the remote endpoint is an RTP mixer or other middlebox
    - The draft mandates support for multiple SSRCs per RTP session, but not for multiple synchronisation contexts (CNAMEs) or for multiple endpoints; should it?
- Do we need to add discussion of SDP signalling for the different scenarios?
    - If so, should it be a separate draft? (JSEP?)

Open Issue: Simulcast

- Broad agreement that simulcast is in scope, but the method for achieving simulcast has to be decided
    - Will be discussed in AVTEXT on Tuesday and MMUSIC on Thursday
- Does simulcast require RTP-level mechanisms beyond those specified?
    - If so, what? draft-westerlund-avtcore-rtp-simulcast-03 is one proposal
    - If not, do we need to specify signalling for simulcast in this draft, or does it go elsewhere? May relate to the W3C API to RTP mapping (later slide…)

Open Issue: Forwarding Media

- Endpoints can participate in multiple RTP sessions
- This potentially lets them forward RTP media data between peers
  - Directly relay RTP packets, acting as an RTP translator
  - Decode then re-encode and transmit the media data
- Should media forwarding be allowed?
  - May be natural to support in the W3C API
  - Requires forwarding browser be aware of congestion state on both paths
  - Two implementation choices exist: browser supports multiple disjoint RTP sessions with media transcoding *or* browser acts as an RTP translator between sessions, forwarding media and translating/forwarding RTCP feedback

Open Issue: Differentiated Services

- Differentiated services possible on a transport flow basis using existing mechanisms
  - Details omitted from this draft – they require no RTP-level mechanisms
  - Sufficient complexity in passing markings between domains, and with the API to mark packets
- Various early proposals to give per-packet marking
  - Use differentiated services field on a per-packet basis
  - Use RTP header extension with deep-packet inspection or middleboxes
  - Proposals are not finished; interaction with congestion control algorithms and AQM is unclear
- Recommendation: this draft outlines the issues, but makes no concrete recommendation

Open Issue: Mapping to W3C API

- The mapping between the W3C API and RTP level concepts has to be agreed and documented
  - Does this go into Section 11 of this draft, or is it part of the W3C API specification?
  - Magnus has a detailed presentation of the issues – propose an ad-hoc discussion meeting later this week to discuss

Open Issue: Correlating Media Streams

- How can we correlate RTP media streams with the signalling? How do we correlate related RTP media streams?
  - Signalled SSRC values or unique payload types per m= line can provide static correlation between SDP m= lines and RTP media flows
    - Limited functionality, but the mechanisms exist to do this already
  - Section 5.2.4: do we need to mandate an RTP header extension that can be used for dynamic correlation of RTP media streams with signalling?
    - RTCP SDES SRCNAME (draft-westerlund-avtext-rtcp-sdes-srcname-03) with RTP header extension for RTCP SDES (draft-westerlund-avtext-sdes-hdr-ext-01) – discuss in AVTEXT
    - Application ID header extension & RTCP SDES item (draft-even-mmusic-application-token-01) – was discussed in MMUSIC this morning
    - Media stream ID (draft-ietf-mmusic-msid-01)
    - May depend on details of the mapping between W3C API and RTP
  - Section 12.2.4: does this draft need to say anything about the signalling for the unified plan? If so, what?

Next Steps

- Resolve these open issues – feedback is needed!

- Submit updated draft and go to WG last call