University of Glasgow | School of Computing Science

Multiple Media Types in an RTP Session

draft-ietf-avtcore-multi-media-rtp-session-01

Magnus Westerlund – Ericsson
Colin Perkins – University of Glasgow
Jonathan Lennox – Vidyo

Status

- Desire to multiplex audio and video in a single RTP session to save ports and ease NAT traversal
  - This draft updates RFCs 3550 and 3551 to allow such multiplexing in certain cases where it is safe to do so, to support WebRTC use cases
- Agreement to adopt this as WG draft at IETF 84
  - The -00 draft was unchanged from the individual submission
  - The -01 draft includes some clarifications, and notes some open issues

Changes in -01

- The main technical changes in -01 are:
  - Update Section 6.1 to overrule text in RFC 3551 that mandates that audio and video are run on separate RTP sessions
  - Update Section 3.3 to clarify what is meant by architectural equality
  - Update Sections 3.3, 5.1, and 6.4 to discuss the constraints imposed by the RFC 3550 requirements for a single RTCP reporting interval
  - Update Sections 6.3 and 7.2 to clarify use with the RTP payload format for Generic FEC
- There are also a number of minor editorial fixes

Single RTCP Reporting Interval Constraint

- RTCP reporting interval derived from:
  - nominal session bandwidth,
  - number of participants and fraction of senders
  - average RTCP packet size
- Base reporting interval, T_d, determines participant timeout period
- Reporting interval for an SSRC does not depend on sending rate of that SSRC
  - Implicit assumption: nominal session bandwidth is representative of all senders in a session
  - All participants have same base reporting interval (roughly: if only a small fraction are senders, they get a slightly smaller base reporting interval)
- This constraint is inherent in RTP design

```
double rtcp_interval(int members,
                     int senders,
                     double rtcp_bw,
                     int we_sent,
                     double avg_rtcp_size,
                     int initial)
{
  double RTCP_SENDER_BW_FRACTION = 0.25;
  double RTCP_RCVR_BW_FRACTION   = 0.75;
  double COMPENSATION            = 2.71828 - 1.5;
  double rtcp_min_time           = initial?2.5:5.0;
  int    n                       = members

  if (senders <= members * RTCP_SENDER_BW_FRACTION) {
    if (we_sent) {
      rtcp_bw *= RTCP_SENDER_BW_FRACTION;
      n = senders;
    } else {
      rtcp_bw *= RTCP_RCVR_BW_FRACTION;
      n -= senders;
    }
  }
  double Td = avg_rtcp_size * n / rtcp_bw;
  if (Td < rtcp_min_time) Td = rtcp_min_time;
  return Td * (drand48() + 0.5)/ COMPENSATION;
}
```

Issue with Reporting Intervals

- Implication of RTCP timing rules: reporting interval can be smaller than desirable for low-rate media in a session with high nominal bandwidth
  - E.g., in a session with high-rate video and low-rate audio, the RTCP reporting interval might be smaller than the audio inter-packet interval, if using the reduced RTCP minimum interval
  - Issue not unique to multiplexed audio & video, but very visible in that case
- Open questions:
  - Is this a problem?
  - If so, do we need to do anything other than give recommended operating regions to avoid the issue? E.g., there might be ways of allowing rapid event reporting with slower regular reports using RTP/AVPF and trr-int – to be investigated further

Issue with SSRC Timeout

- Need to ensure consistent RTCP timeout interval
  - Either all participants use 5 second T_{min} for base RTCP reporting interval, *or* all use the reduced minimum T_{min} – cannot mix-and-match, else get spurious timeouts
- RFC 3550 suggests using the 5 second minimum, but is inconsistent – need to clarify

Next Steps

- Where should we address these issues? In this draft, or in a separate clarification to RFC 3550
  - The issues are highlighted by multiplexed audio and video, but are also present in other scenarios